



FACULTY OF BUSINESS AND ECONOMICS

DISSERTATION

---

Political behavior in the social media age:  
a data mining approach

---

Thesis submitted for the degree of Doctor of Applied Economics  
at the University of Antwerp to be defended by  
Stiene PRAET

*Author:*  
Stiene PRAET

*Supervisors:*  
Prof. dr. ir. David MARTENS  
Prof. dr. Peter VAN AELST

Antwerp, 2021

**Supervisors**

Prof. dr. ir. David MARTENS (University of Antwerp)

Prof. dr. Peter VAN AELST (University of Antwerp)

**Members of the Examination Committee**

Prof. dr. Kenneth SÖRENSEN (University of Antwerp)

Prof. dr. ir. Wouter VERBEKE (KU Leuven)

Prof. dr. Silke ADAM (Universität Bern)

Prof. dr. Fosca GIANNOTTI (ISTI-CNR Pisa)

© Copyright 2021, Stiene PRAET

This research was supported by the University of Antwerp [FFB160362 BOF GOA].

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without permission in writing from the author.

## Publications and conference presentations

---

### Publications

The following publications are included in part or in an extended version in this thesis:

- Praet, S., Van Aelst, P., Daelemans, W., Walgrave, S., Kreutz, T., Peeters, J., and Martens, D. (2020a). Comparing automatic content analysis methods to distinguish issue communication by political parties on Twitter. *Accepted: Computational Communication Research*.
- Praet, S., Van Aelst, P., Van erkel, P., Van der Veecken, S., and Martens, D. (2020c). How Facebook likes can add lifestyle politics to complement traditional election survey research. *R&R: EPJ Data Science*.
- Praet, S., Guess, A. M., Tucker, J. A., Bonneau, R., and Nagler, J. (2021a). What's Not to Like? Facebook Page Likes Reveal Limited Polarization in Lifestyle Preferences. *R&R: Political Communication*.
- Praet, S., Van Aelst, P., and Martens, D. (2021b). Patterns of democracy? Social network analysis of parliamentary Twitter networks in 12 countries. *R&R: Online Social Networks and Media*.

Furthermore, the following publications were part of my PhD research, but are not covered in this thesis. The topics of these publications are outside of the scope of the material covered here:

- Praet, S. and Martens, D. (2020). Efficient parcel delivery by predicting customers' locations. *Decision Sciences*, 51(5):1202–1231.
- Peeters, J., Van Aelst, P., and Praet, S. (2019). Party ownership or individual specialization? A comparison of politicians' individual issue attention across three different agendas. *Party Politics*, pages 1–12.
- Stankova, M., Praet, S., Martens, D., and Provost, F. (2021). Node classification over bipartite graphs through projection. *Machine Learning*, 110:37–87.

- Ramon, Y., Martens, D., Evgeniou, T., and Praet, S. (2021). Metafeatures-based rule-extraction for classifiers on behavioral and textual data. *Accepted: Machine Learning*.
- Luhiste, M., Popa, S. A., Theocharis, Y., Barberá, P., Praet, S., Fazekas, Z., and Tucker, J. (2021). When does Abuse and Harassment Marginalize Female Political Voices on Social Media? *To be submitted*.

### Conference presentations

- Praet, S. and Martens, D. (2017). Efficient Parcel Delivery by Predicting Customers' Locations. IFORS: 21st Conference of the International Federation of Operational Research Societies, Quebec, Canada.
- Praet, S., Van Aelst, P., Daelemans, W., Walgrave, S., Kreutz, T., , and Martens, D. (2018b). Issue Ownership and Communication by Political Parties on Twitter. EURO: 29th European Conference on Operational Research, Valencia, Spain.
- Praet, S., Daelemans, W., Kreutz, T., Van Aelst, P., Walgrave, S., and Martens, D. (2018a). Issue communication by political parties on Twitter. Data Science, Journalism and Media workshop at KDD Conference, London, UK.
- Praet, S., Van Aelst, P., Van der Veekend, S., and Martens, D. (2019b). What I Do versus Who I Am. Comparing Facebook likes with survey data to gain insights in political preference. EURO: 30th European Conference on Operational Research, Dublin, Ireland.
- Praet, S., Van Aelst, P., Van der Veekend, S., and Martens, D. (2019a). Predictive modeling on Facebook likes to gain insights in voter profiles in a multi-party system. 5th International Conference on Computational Social Science IC2S2, Amsterdam, The Netherlands.
- Praet, S., Van Aelst, P., and Martens, D. (2020b). Partisanship in everyday life: predictive modeling with Facebook like data. ORBEL34: 34th conference of the Belgian Operational Research Society, Lille, France.

## Acknowledgements

---

The photograph on the cover of this thesis was shot<sup>1</sup> during my trip to Washington D.C. in November 2019. It's hard to ignore the 'Second Amendment Rally' as a precursor to the events that would take place at this same building one year later in January 2021. The storming of the Capitol is an illustration of the potential negative impact of social media on our political lives. Social media is an exciting technology that has connected people worldwide, but has also led to increased division and polarization. A very intriguing and topical subject, and I am forever grateful for the opportunity to do research in this area, for all the inspiring people I've met and worked with, and for everything I learned from them.

First and foremost, I would like to express my immense gratitude to my supervisors David Martens and Peter Van Aelst. David, thank you for believing in me and introducing me to your research passion. I appreciate the trust you gave me so that I could find my own way in the research field, while at the same time guiding and supporting me when needed. Next to that, I got to know you as a fun and positive person who is always in for dinner, drinks, or an escape room/zombie shooting. Peter, thank you for welcoming me to "de GOA". You have shown me that political research can be fun (when adding big data to it of course). Our collaboration has encouraged me to step beyond my comfort zone and explore a new research field, which made my PhD journey all the more enriching. Thank you both for the time and effort you invested in this research, for the great conversations and advice that have helped me grow both professionally and personally, and for the many additional opportunities you gave me, including presenting my work at international conferences and an unforgettable research stay at New York University. I wish you both the best in your future endeavours, professionally as well as personally.

I would like to thank the members of my Examination Committee, Kenneth Sörensen, Wouter Verbeke, Silke Adam, and Fosca Giannotti, for their time and thorough evaluation of my work. Your valuable questions, comments, and feedback have helped me to improve the quality of this thesis in the final stage of my PhD trajectory. Thank you Nele, Joeri, Kelly, and Kim for the organization and technical support of my online PhD defence.

Furthermore, I would like to thank all other GOA (NWS Data) members, Stefaan, Walter, Steve, Patrick, Kenza, Jeroen, and Tim, for the productive collaboration and fun meetings (especially when there was food). I hope we can finish the project in

---

<sup>1</sup> Credits to Benjamin Bergers (@truth.in.motion)

style with one more fantastic closing activity to come. I would like to thank Kenza, Jeroen, and Tim in particular for embarking on this PhD journey together with me, for our frequent coffee dates, and for the –well-intentioned– research collaborations. It was a pleasure to have shared this experience with you and I wish you all the best on finishing your PhDs and beyond.

I am grateful for the Amazing Data Mining team. Thank you Marija, Ellen, Sofie, Jellis, Yanou, Tom, Dieter, Raphael, and Sofie (G.), for the technical support, interesting discussions, useful advice, productive collaborations, conference trips, fun afterworks and so much more. As a proud senior I now pass on the honour to Yanou. I am sure you will do an excellent job.

These five years wouldn't have been as enjoyable without the support and friendship of my amazing office mates. Thank you Ellen and Sofie, for welcoming me so warmly into your office on my very first day at university. I have loved sharing the office, cat and dog stories, and a trip to London! Thank you Steven, Jellis, Yanou, and Dieter for the fun vibe at our office which attracted people from all over the fifth floor to relax in our chair, talk about ping-pong or Loenhout parties while enjoying a freshly-brewed coffee. Sofie (G.), your cheerful presence at group runs, sport sessions, and other social activities greatly compensates for the lack of time we could spend together at the office. Thank you for being such a dedicated Secret Santa and colleague.

I thank all my wonderful colleagues at the department and faculty for the unforgettable coffee breaks, lunches, evenings at Barracuda, weekends, departmental days, group runs, conferences and many more. Thank you: Jasmine, for founding the ENM running team, which has become a valuable social activity during times of restrictions. Lissa, for being the most loyal (running) buddy, and for initiating so many fun activities. Thank you for the beautiful drawing below. Babiche, for our first day at work, our first conference in Quebec, and for the many more trips that followed. Floor, for your happy vibes and caring personality, which we all miss a lot. Yanou, for visiting me in New York and for your hospitality during fun evenings. Kelly, for the pleasant company during runs and taekwondo sessions. Tom, for having a witty comment at hand in every situation. Loïc and Wito, for your amusing visits to the fifth floor. Amy, for the fun times at Orbel, and for the plant advice. Tyché, for being the kindest person on earth. Florian, for your great research advice, and your talent to get involved in comic situations. Sophie and Carlos, for our trip in Washington D.C. Paul and Renata, for the extensive –late night– discussions. Young Research Network and the FBE weekendcommittee, for bringing us together. Kevin, for your engagement with the department and the boardgames. Mauricio, for being everybody's best friend. Trijntje and Johan S., for the in-depth debates during coffee breaks. Kenneth, for the funny anecdotes. Maxim, Iolanda, Gwenny, Hanne, Sam, Adriaan, Heidi, Johan B., Ying, Michell, Bryan, Cristina, An, Aline, Lenore, ... Thank you for being or having been the greatest colleagues.


My gratitude to Joshua Tucker for having me at the CSMAP research group at New York University. Thank you Andy Guess, Joshua Tucker, Richard Bonneau and Jonathan Nagler for the opportunity to collaborate on Chapter 6 of this thesis. Many thanks to the brilliant CSMAP team for the interesting presentations and great conversations but mostly for the lunches, drinks, and recommendations around NYC. Thank you to my new friends for making this stay unforgettable and continuing our friendship at home.

Over the years I have been blessed with the greatest friends, who provided moral support and welcome distractions. To my dear friends, thank you for the unforgettable memories, the laughter, and for always having my back. Stay awesome!

To my parents, thank you for your unconditional love and support and for all the chances you gave me. You have always been and still are my role models. To my –extended– family, you are the friends I didn’t choose. Thank you for bringing me so much warmth and joy. I am looking forward to meeting the latest extension to our family in July.

And finally, from (too) rigorously reviewing my writing, to pimping my conference presentations, and cheering me up after a hard day’s night. Thank you Benjamin for your endless support, I am looking forward to our new adventures together.

Thank you!  
Stay connected.



Stiene Praet  
Antwerp, June 2021





## Abstract

---

The storming of the United States Capitol in January 2021 dramatically illustrates the impact of social media on society and political outcomes. The new reciprocal relationships afforded by digital media have reshaped the way political information is produced and consumed, and challenge some of the established theoretical insights in political communication. At the same time, the digital revolution also offers new opportunities for empirical research. By leveraging the information captured in digital traces, we can expand our understanding of political behavior in a way that was simply unimaginable a mere decade ago.

In this PhD thesis, I collect and analyze social media data to explore the opportunities of data mining, text mining, and network analysis techniques for political research. The first part studies elite polarization with a large-scale comparison of political Twitter networks in 12 countries. In addition, a more in-depth study of party communication in Belgium is performed. In the second part, I analyze political polarization in non-political domains using Facebook-like-data in Belgium and the United States. I find that political polarization and partisanship are dependent upon the social network, institutional context, and individual characteristics. To mitigate polarization, I suggest lifestyle domains in which most cross-cutting interactions are present.

This thesis shows that social media data provide a unique and rich source of online behavior but also come with ethical, technical, and methodological challenges. Therefore, joint efforts between social and computer scientists are needed to convert the enormous empirical potential into valuable insights.



# Dutch Abstract

---

## POLITIEK GEDRAG OP SOCIALE MEDIA: EEN DATAMINING TOEPASSING

De bestorming van het Amerikaanse Capitool in januari 2021 illustreert op dramatische wijze de impact van sociale media op onze samenleving en politiek. De nieuwe wederkerige relaties die mogelijk worden gemaakt door digitale media hebben de manier waarop politieke informatie wordt geproduceerd en geconsumeerd drastisch hervormd. Zo worden enkele van de gevestigde theoretische inzichten in politieke communicatie opnieuw in vraag gesteld. Tegelijkertijd biedt de digitale revolutie ook nieuwe kansen voor empirisch onderzoek. De digitale voetsporen die we allemaal achterlaten op het internet kunnen wetenschappers helpen om politiek gedrag beter te begrijpen.

In dit proefschrift verzamel en analyseer ik sociale media data om de mogelijkheden van datamining, textmining en netwerkanalysetechnieken voor politiek onderzoek te verkennen. Het eerste deel bestudeert elitepolarisatie met een grootschalige vergelijking van politieke Twitter-netwerken in 12 landen. Daarnaast wordt er een meer diepgaande studie van de partijcommunicatie in België uitgevoerd. In het tweede deel analyseer ik politieke polarisatie in niet-politieke domeinen met behulp van Facebook-like-data in België en de Verenigde Staten. Mijn bevindingen leren dat politieke polarisatie en partijdigheid afhankelijk zijn van het sociale netwerk, de institutionele context en individuele kenmerken. Om polarisatie te verminderen, stel ik interessegebieden voor die burgers met verschillende politieke overtuigingen met elkaar kunnen verbinden.

Dit proefschrift toont aan dat gegevens van sociale media een unieke en rijke bron van online gedrag vormen, maar ook ethische, technische en methodologische uitdagingen met zich meebrengen. Daarom zijn gezamenlijke inspanningen van sociale- en computerwetenschappers nodig om het enorme empirische potentieel van dit soort data om te zetten in waardevolle inzichten.



# Contents

---

1	INTRODUCTION	1
1.1	A connected world	2
1.2	Social media and politics	4
1.2.1	A new information flow	5
1.2.2	From liberation to repression	7
1.3	Computational social science	8
1.3.1	Two research paths	8
1.3.2	Opportunities for empirical research	12
1.3.3	Challenges	13
1.3.4	Privacy and ethical concerns	15
1.4	Social media to study political behavior	18
1.5	Contribution and framework	20
2	DATA AND METHODS	25
2.1	Data collection	26
2.1.1	Twitter data (Part I)	26
2.1.2	Facebook data (Part II)	27
2.2	Data mining concepts	31
2.2.1	Classification techniques	31
2.2.2	Dimensionality reduction	34
2.2.3	Performance metrics	35
i	TWITTER TO STUDY COMMUNICATION AND INTERACTIONS	39
3	ISSUE COMMUNICATION ON TWITTER	41
3.1	Introduction	42
3.2	Automated content analysis	44
3.3	Data and methods	47
3.3.1	Data collection	47
3.3.2	Preprocessing of tweets	49
3.3.3	Tweet representation	49
3.3.4	Classification models	52
3.3.5	Evaluation	52
3.4	Results	53
3.4.1	Comparison of tweet representations	53
3.4.2	Which issues separate party communication from other parties?	55
3.4.3	How consistent is party communication?	57
3.5	Conclusion and future research	57
4	PARLIAMENTARY TWITTER NETWORKS	61
4.1	Introduction	62

## CONTENTS

4.2	Parliamentary Twitter networks and elite polarization . . . . .	63
4.3	Data collection . . . . .	65
4.4	Methods . . . . .	66
4.4.1	Network topology . . . . .	67
4.4.2	Inter-party communication . . . . .	69
4.5	Results . . . . .	70
4.5.1	Twitter usage . . . . .	70
4.5.2	Network topology . . . . .	72
4.5.3	Inter-party communication . . . . .	81
4.6	Conclusion . . . . .	82
ii	FACEBOOK LIKES TO STUDY LIFESTYLE AND POLARIZATION . . . . .	85
5	FACEBOOK LIKES TO STUDY LIFESTYLE POLITICS . . . . .	87
5.1	Introduction . . . . .	88
5.2	Predictive modeling with Facebook Likes . . . . .	90
5.3	Data collection . . . . .	91
5.4	Methods . . . . .	94
5.5	Results . . . . .	97
5.5.1	Predictive performance . . . . .	97
5.5.2	Insights . . . . .	98
5.6	Conclusion and future research . . . . .	101
6	POLARIZATION IN LIFESTYLE DOMAINS . . . . .	105
6.1	Introduction . . . . .	106
6.2	Polarization, lifestyle, and social media . . . . .	107
6.3	Hypothesis and research questions . . . . .	108
6.4	Data collection . . . . .	110
6.5	Methods . . . . .	113
6.6	Results . . . . .	116
6.6.1	Facebook categories . . . . .	119
6.6.2	User-level analysis . . . . .	121
6.7	Conclusion . . . . .	126
7	POLARIZATION IN BELGIUM COMPARED TO THE U.S. . . . .	129
7.1	Introduction . . . . .	130
7.2	Data and methods . . . . .	130
7.3	Results . . . . .	131
7.3.1	Facebook categories . . . . .	131
7.3.2	User-level analysis . . . . .	134
7.4	Conclusion . . . . .	141
iii	CONCLUSIONS . . . . .	143
8	CONCLUSIONS . . . . .	145
8.1	Main findings . . . . .	146
8.2	The future of computational social media research . . . . .	149
8.3	The future of social media . . . . .	153

iv	APPENDICES	177
9	APPENDICES	179
9.1	Issue communication on Twitter . . . . .	179
9.1.1	Evaluation of the CAP dictionary per issue . . . . .	179
9.1.2	Extended dictionary with word embedding . . . . .	182
9.1.3	Most discriminative features for the data-driven representations	183
9.2	Parliamentary Twitter networks . . . . .	185
9.2.1	Data collection . . . . .	185
9.2.2	Visualization of parliamentary Twitter networks . . . . .	188
9.2.3	Community detection . . . . .	189
9.2.4	Inter-party communication and ideology . . . . .	190
9.3	Facebook likes to study lifestyle politics . . . . .	191
9.3.1	Data collection . . . . .	191
9.3.2	Association methods . . . . .	194
9.3.3	Statistical comparison of models . . . . .	199
9.3.4	Voter profiles . . . . .	201
9.4	Polarization in lifestyle domains . . . . .	205
9.4.1	Sample details . . . . .	205
9.4.2	Coding categories . . . . .	206
9.4.3	Methods and measures . . . . .	209
9.4.4	Robustness checks . . . . .	209
9.4.5	Results . . . . .	217
9.5	Polarization in Belgium compared to the U.S. . . . .	223

## List of Figures

---

Figure 1.1	Worldwide social media use . . . . .	3
Figure 1.2	Traditional and new model of political information flow . . . .	6
Figure 1.3	Traditional versus fine-grained behavioral data . . . . .	10
Figure 1.4	Exploratory research methodology . . . . .	20
Figure 1.5	Contributions of this PhD thesis . . . . .	23
Figure 2.1	Ethical considerations Twitter . . . . .	27
Figure 2.2	Facebook Login button . . . . .	28
Figure 2.3	Webpage for data collection . . . . .	29
Figure 2.4	Ethical considerations Facebook . . . . .	30
Figure 2.5	Example of classification model . . . . .	31
Figure 2.6	Logistic function . . . . .	32
Figure 2.7	Illustration of matrix factorization . . . . .	34
Figure 2.8	Illustration of ROC curve . . . . .	37
Figure 3.1	Schematic representation of methods . . . . .	47
Figure 3.2	Overview of our exploratory approach . . . . .	48
Figure 3.3	Comparison of our three methods . . . . .	54
Figure 4.1	Plot of ideological distance and inter-party relations . . . . .	70
Figure 4.2	Proportion of politicians with Twitter account . . . . .	71
Figure 4.3	Average and median number of tweets . . . . .	72
Figure 4.4	Average and median number of followers . . . . .	72
Figure 4.5	Parliamentary networks in The Netherlands . . . . .	74
Figure 4.6	Hierarchical clustering . . . . .	77
Figure 4.7	Network visualization for the retweets networks . . . . .	79
Figure 4.7	Network visualization for the retweets networks (Continued)	80
Figure 5.1	Average AUC for the Facebook categories . . . . .	98
Figure 5.2	Insights into voter profiles . . . . .	100
Figure 6.1	Page Likes distribution . . . . .	112
Figure 6.2	Bipartite network Like relations . . . . .	113
Figure 6.3	Network visualization and modularity . . . . .	117
Figure 6.4	Page ideology distribution . . . . .	118
Figure 6.5	Homogeneity distribution . . . . .	123
Figure 7.1	Page ideology distribution . . . . .	132
Figure 7.2	Facebook pages per category . . . . .	135
Figure 7.2	Facebook pages per category (Continued) . . . . .	136
Figure 7.2	Facebook pages per category (Continued) . . . . .	137
Figure 7.3	Homogeneity distribution . . . . .	138
Figure 9.1	The parliamentary mentions network at party level. . . . .	188



Figure 9.2	The parliamentary retweets network at party level. . . . .	189
Figure 9.3	Coverage association methods . . . . .	199
Figure 9.4	Page ideology distribution for all pages . . . . .	210
Figure 9.5	Homogeneity distribution for all pages . . . . .	210
Figure 9.6	Network visualization bottom nodes . . . . .	218
Figure 9.7	Facebook pages per category . . . . .	220
Figure 9.7	Facebook pages per category (Continued) . . . . .	221
Figure 9.7	Facebook pages per category (Continued) . . . . .	222
Figure 9.8	Facebook users . . . . .	223

## List of Tables

---

Table 1.1	Social science versus computer science. . . . .	14
Table 1.2	Examples of research applications in social media and politics	21
Table 2.1	Confusion matrix . . . . .	36
Table 3.1	Traditional dictionary approach . . . . .	46
Table 3.2	The number of accounts and tweets per party . . . . .	48
Table 3.3	Overview of the 21 CAP issues (Sevenans et al., 2014). . . . .	50
Table 3.4	Most discriminative features Vlaams Belang (BoW) . . . . .	54
Table 3.5	Most discriminative features Vlaams Belang (topic model) . . . . .	55
Table 3.6	Issue communication of Flemish parties . . . . .	56
Table 3.7	Classification performance and interpretability . . . . .	57
Table 4.1	Overview of the countries in our study. . . . .	66
Table 4.2	Ideological distance and inter-party follower relations . . . . .	70
Table 4.4	number of politicians with a Twitter account . . . . .	71
Table 4.5	Description of the followers networks . . . . .	74
Table 4.6	Description of the mentions networks . . . . .	75
Table 4.7	Description of the retweets networks . . . . .	75
Table 4.8	E-I index per country . . . . .	81
Table 4.9	Correlation inter-party relations and ideological distance . . . . .	82
Table 5.1	Facebook categories . . . . .	93
Table 5.2	Datasets and variables . . . . .	96
Table 5.3	Predictive performance . . . . .	98
Table 6.1	Sample details . . . . .	111
Table 6.2	Facebook categories . . . . .	112
Table 6.3	Homogeneity and ideology for Facebook categories . . . . .	119
Table 6.4	Average number of Page Likes per ideology . . . . .	121
Table 6.5	Examples of pages at percentiles . . . . .	121
Table 6.6	Proportion of users with cross-cutting interactions . . . . .	122
Table 6.7	Average number of likes . . . . .	123
Table 6.8	Determinants of individual homogeneity per category . . . . .	124
Table 7.1	Comparison of homogeneity per category . . . . .	133
Table 7.2	Proportion of respondents with individual homogeneity score higher than 0.1 . . . . .	134
Table 7.3	Determinants of individual homogeneity with ideology . . . . .	139
Table 7.4	Determinants of individual homogeneity with party preference	140
Table 9.1	Intercoder reliability for the annotated tweets . . . . .	179
Table 9.2	Evaluation of the CAP dictionary . . . . .	180
Table 9.3	Issues assigned using CAP dictionary . . . . .	181
Table 9.4	Most discriminative BoW . . . . .	183

Table 9.5	Most discriminative topics . . . . .	184
Table 9.6	Number of politicians and sources per country. . . . .	185
Table 9.7	Detected communities and adjusted mutual information . . .	189
Table 9.8	RILE scores . . . . .	190
Table 9.10	Sample details . . . . .	193
Table 9.11	Gender, age and education of the Belgian population . . . . .	194
Table 9.12	Example contingency table . . . . .	194
Table 9.13	Expected values for left ideology . . . . .	196
Table 9.14	Comparison of the methods and their characteristics. . . . .	198
Table 9.15	Comparison of prediction models for ideological leaning . . .	199
Table 9.16	Comparison of prediction models for party preference . . . .	200
Table 9.17	Most discriminative pages for ideological leaning . . . . .	201
Table 9.18	Most discriminative survey questions for ideological leaning .	202
Table 9.19	Most related NMF components to left ideological leaning . .	203
Table 9.20	Most related NMF components to right ideological leaning . .	204
Table 9.21	Comparison of samples. . . . .	205
Table 9.22	Total number of Page Likes per ideology . . . . .	206
Table 9.23	Categories codebook . . . . .	208
Table 9.24	Entropy, Cramer's V, and Variance . . . . .	212
Table 9.25	Determinants of individual homogeneity for different homo- geneity metrics. . . . .	213
Table 9.26	Determinants of individual homogeneity when regression is weighted for individual number of page likes . . . . .	215
Table 9.27	Determinants of individual homogeneity with party identifi- cation (PID) strength . . . . .	216
Table 9.28	Communities bottom node projection . . . . .	217
Table 9.29	Communities top node projection . . . . .	219
Table 9.30	Overlapping coefficients U.S. . . . .	223
Table 9.31	Overlapping coefficients Belgium . . . . .	223



## Introduction

---

*“Technology is neither good nor bad, nor is it neutral.”*

*Kranzberg’s First Law of Technology*

In 2010, Mark Zuckerberg was named Time Magazine’s Person of the Year, “for connecting more than half a billion people and mapping the social relations among them. For creating a new system of exchanging information and for changing how we live our lives” (Time, 2010). A mere eight years later, the same man testified before American congress, in the aftermath of the platform’s misuse by the data analytics firm Cambridge Analytica. It’s unquestionable that the rise of social media has revolutionized our modern society and has altered our social, economic and political behavior. Yet, the dangers of these highly unregulated, non-transparent and algorithm-driven online platforms are just around the corner: addiction, privacy breaches, misinformation, polarization, hate speech and downright manipulation of our democratic system (Tucker et al., 2017). Social media has not only changed how politics is practiced, but also how it is being studied. The digital revolution calls for innovative methods to benefit from the massive amounts of digital data that have come available.

In this thesis I explore the opportunities of data mining, text mining and network analysis techniques to study political behavior and communication on a fine-grained level. The first part studies online political communication with a large-scale comparison of political Twitter networks in 12 countries, and a more in-depth study of issue communication in Belgium. In the second part I analyze political polarization in non-political domains using Facebook Like data in Belgium and the United States. This introduction chapter outlines the current social media landscape, and the influence of social media on the political game and our democracy. Next, I discuss how social media data can be used to study online political behavior, and the ethical, technical and methodological challenges that come with social media research. Finally, the contributions of this PhD thesis to the field are highlighted.

### 1.1 A connected world

Thanks to social media, protesters in the Middle East could communicate and organize themselves during the Arab Spring (Howard et al., 2011), phenomenal funds were raised and people were kept informed about the Australian bush fires in the summer of 2019 (Mack, Graig, 2020), and we could all keep our loved ones close in times of physical distancing during the corona pandemic. Social media connects people across the world, reinforces friendships, fosters interaction, provides access to information, and gives a voice to the wider public. Social media arguably has transformed every aspect of our social lives and is now one of the most prevailing (digital) activities worldwide.

Since the launch of Facebook in 2004, social media platforms have exploded in popularity. The number of global social media users (Figure 1.1a) has risen from 2.86 billion since the start of this PhD research in 2017 to 3.6 billion users by the end of 2020, and is projected to continue its steady growth in the future (Clement, 2020c). A significant part of our day, on average 2.5 hours,<sup>1</sup> is spent on social media (see Figure 1.1c), and among the most popular social media platforms are Facebook, Youtube and WhatsApp (Figure 1.1b). Social media is mainly used to follow people, to like and comment on posts, to post and share content, and to send private messages (Kunst, 2020).

Although most of us could easily list a plethora of social media platforms, it is more challenging to agree on a formal definition of what social media is. Generally, social media can be defined as "a computer-based technology that facilitates the sharing of ideas, thoughts, and information through the building of virtual networks and communities" (Dollarhide, 2020, para. 1). Within this general definition however, various types of social media can be distinguished further: social networking sites (e.g., Facebook), blogs and applications such as collaborative projects (e.g., Wikipedia), content communities (e.g., YouTube), and virtual games and social worlds (e.g., World of Warcraft, Second Life) (Kaplan and Haenlein, 2010). In this research, we will focus on the first type, social networking sites:

*web-based services that allow individuals to (1) construct a public or semi-public profile within a bounded system, (2) articulate a list of other users with whom they share a connection, and (3) view and traverse their list of connections and those made by others within the system* (Boyd and Ellison, 2007, p.211).

More specifically, in the first part of this PhD thesis we will focus on the social networking site Twitter. Twitter is a free microblogging service that allows users

<sup>1</sup> The Belgian average is slightly lower: people spend 1.5 hours per day on social media (We Are Social and Hootsuite, 2020).

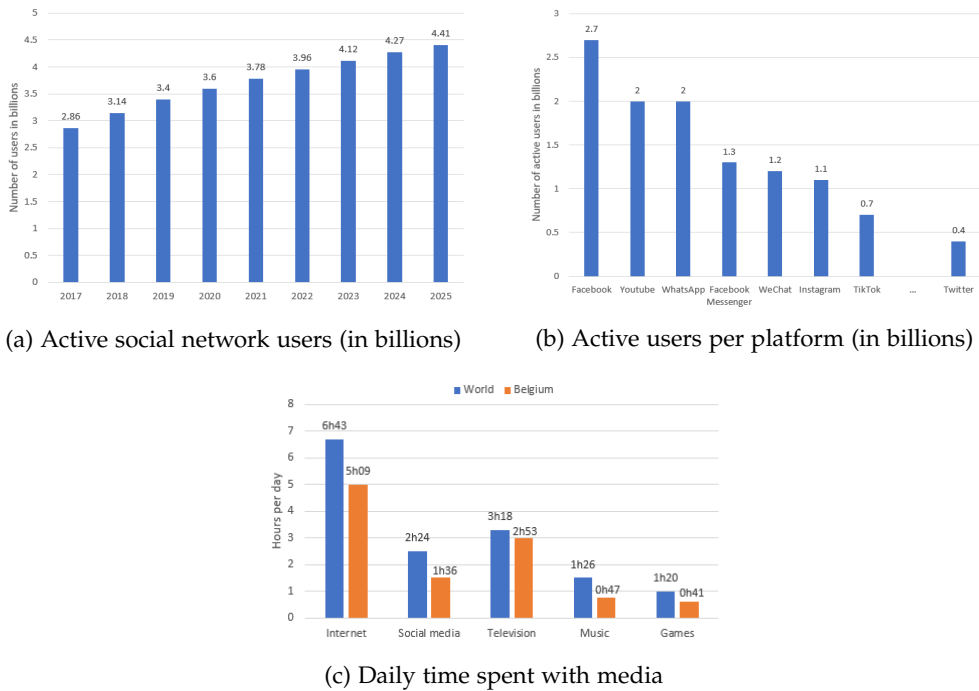


Figure 1.1: Number of social network users worldwide from 2017 to 2025 (Clement, 2020c) (a), most popular social networks worldwide as of October 2020 (Clement, 2020b) (b), and daily time spent with media (We Are Social and Hootsuite, 2020) (c)

to post short messages of 280 characters.<sup>2</sup> Globally, the platform reaches nearly 400 million people, or almost 6% of the population older than 13. More than half of the Twitter audience is between 18 and 34 years old and the ratio of female to male Twitter users is roughly two to three (We Are Social and Hootsuite, 2020). In the United States, Twitter is especially popular in urban areas, among those with high education and income (Perrin and Anderson, 2019). Most users rarely post “tweets” themselves, and the most active 10% of Twitter accounts create 80% of tweets (Wojcik and Hughes, 2019). Twitter reaches 1.3 million people in Belgium, which is 13% of the population aged 13 and older. Belgian Twitter users are younger and roughly 70% are men. Twitter is especially popular among politicians and political parties to communicate with the public (Jungherr, 2016; Vargo et al., 2014). U.S. congress is very active on the platform (van Kessel et al., 2020), and more than three quarters of the members of the European Parliament have installed a Twitter account (Scherpereel et al., 2017).

<sup>2</sup> Originally, the character limit was 140 but by the end of 2017 Twitter decided to expand the character limit to 280, see [https://blog.twitter.com/official/en\\_us/topics/product/2017/Giving-you-more-characters-to-express-yourself.html](https://blog.twitter.com/official/en_us/topics/product/2017/Giving-you-more-characters-to-express-yourself.html).

In the second part of the thesis we will analyze data from the social networking platform Facebook. Currently, Facebook is the most popular social network worldwide (Clement, 2020b) and is widely considered as one of the Big Four tech companies, along with Google, Apple, and Amazon (all together known under the acronym GAFA). The company also owns the mobile messaging apps WhatsApp and Facebook Messenger, as well as photo-sharing app Instagram. Facebook users worldwide are predominantly male and between the age of 18 to 44 years (Clement, 2020a). Although younger age groups are over-represented, Pew Research Center found that Facebook use in the United States is also relatively common across older age groups. According to their findings, almost 70% of people aged 50 to 64 and nearly half of those 65 and older said they use the site. While the use of Facebook is relatively evenly spread, the use among low-educated and low-income groups is slightly less. The majority of Facebook users visit the site daily, or even several times a day (Perrin and Anderson, 2019). In Belgium, Facebook is by far the most popular social media platform, with almost 66% of the Belgian population using the platform. Most Facebook users in Belgium are between 25 and 34 years old, and, in contrast to the worldwide numbers, more often female than male (Statista, 2020).

We can no longer imagine a world without social media. However, after many years of frequent social media use, we have to recognize that the promising technology also has its drawbacks. Well-developed strategies and algorithms exploit our social nature to create a “Hype Machine”, designed to keep our attention and manipulate our actions (Aral, 2020). Infinite scrolling and push notifications trigger addiction and mental health problems, our personal data can be misused to influence our behavior, and efficient algorithms and misinformation drive us into filter bubbles, polarization, and extremism. Or, in the words of David Carroll: “How did the dream of the connected world tear us apart?” (Ironically, he posted this on Twitter.<sup>3</sup>)

## 1.2 Social media and politics

Social media is not only omnipresent in our personal lives, the presidential term of Donald Trump in the United States leaves no doubt about the importance of social media for society and political outcomes. His term started with a controversial social media campaign and ended with inflammatory calls on social media not to accept the democratic election results,<sup>4</sup> after which his account was eventually banned from several platforms (CNN, 2021b).

<sup>3</sup> David Carroll, Associate Professor of Media Design at Parson, tweeted this in reaction to a story about a gunman’s post on Instagram. He featured in the Netflix documentary *The Great Hack* (2019) and is known for challenging Cambridge Analytica and the SCL Group companies in the UK High Court.

<sup>4</sup> Arguably, leading to an assault of Trump supporters on the U.S. Capitol in Washington D.C on January 6, 2021 (CNN, 2021c), which in turn led to a historic second impeachment of Donald Trump. The impeachment trial ended once again with an acquittal on February 13, 2021 (CNN, 2021a).



Donald Trump's tremendous fame, huge following, and very great skill<sup>5</sup> in navigating the digital landscape gave him a unique asset in terms of social media dominance. In 2016, he had one third more followers than his opponent Hillary Clinton, his tweets were retweeted more than three times as often, and his Facebook posts were re-shared five times as much (Pew Research Center, 2016). On top of that, the Trump campaign spent more on social media than the Clinton campaign and worked together with the micro-targeting firm Cambridge Analytica (Persily, 2017). Cambridge Analytica had collected Facebook data —unlawfully, as it turned out later. They used this data to develop psychometric profiles and sequentially to micro-target persuadable voters. Even more controversially, they targeted Clinton supporters to reduce voter turnout among them (González, 2017).

The impact of social media reached far beyond the candidates' campaigns however, to being a tool of outside actors to undermine democracy (Tucker et al., 2017). Publishing pro-Trump and anti-Clinton stories proved to be a very profitable business. The more outrageous the stories were, the more visitors they would attract and, in turn, the more clicks on the advertisements appearing on the page. For example, the story about an FBI agent being killed after leaking Clinton's emails was completely false but was shared over half a million times (Sydell, 2016). These false stories did not only sabotage the elections by misinforming the public at large, some induced direct danger and violence. "Pizzagate" was based on conspiracy theories about a child-trafficking ring involving Hillary Clinton and her campaign chairman. A man who believed the theory opened fire with a rifle at a pizza-restaurant in Washington D.C. (Tangherlini et al., 2020). The power of fake news is determined by its speed of dissemination. With the help of paid trolls<sup>6</sup> and automated bots, misinformation can be spread very efficiently. This became painfully clear during the 2016 elections.

Although the scientific community is still trying to understand the direct influence of these events on the outcome of this and other elections (Allcott and Gentzkow, 2017; González, 2017; Rathi, 2019; Boerboom, 2020; Bovet and Makse, 2019), the 2016 U.S. presidential election can be seen as a tipping point from a general belief in the pro-democratic effects of social media towards a more negative view on the challenges that social media poses for democracy (Tucker et al., 2017).

### 1.2.1 *A new information flow*

The social media revolution has altered the political sphere by shifting who controls, consumes and distributes political information. In a democracy, citizens need information about politics in order to hold informed opinions and act meaningfully. Traditionally, the main channels for political communication were newspapers, radio, and television. This represents a top-down model of communication, where political actors and traditional news media are the dominant players in the production of

---

<sup>5</sup> THE best, really, it's true

<sup>6</sup> Paid teams that post on social media to influence public opinion




Figure 1.2: The traditional top-down model of political information (a) and the new reciprocal relationships afforded by social media (b)

political news (Johansson, 2019). Although occasional opinion polls and market research were used to capture what the public thinks, the public remained a rather insignificant player (Figure 1.2a). Social media changes two vital elements of the political information flow (Gainous and Wagner, 2013). First, it enables citizens to self-select information<sup>7</sup>—often in a way that is congruent with their own beliefs and ideas—and actively participate in the production and especially the distribution of content. This creates a so-called “many-to-many” communication structure. Second, it allows politicians and political parties to shape their content and communicate directly with voters without the intervention of traditional media. Social media has given the public a stronger voice and enables new reciprocal relationships in political communication (Figure 1.2b).

Instead of journalistic gate-keepers deciding what information reaches us, the selection is increasingly determined by algorithmic optimizations (Sirbu et al., 2019). The information flow in this new model is mediated through what Van Dijck and Poell (2013) call *social media logic*. Social media logic refers to the processes through which these platforms channel social traffic and includes four basic elements: programmability, popularity, connectivity, and datafication. First, programmability are the technological mechanisms (e.g. computer code, data, algorithms, protocols, and interfaces) that influence users’ experiences on consuming and creating information. In turn, users retain significant agency in the process of steering information creation through their own contributions and interactions with the platform mechanisms. Second, popularity refers to the mechanisms (both algorithmic and socio-economic) to boost popularity of people or content. In spite of the platforms’ egalitarian promise, in a sense that all users can equally participate and contribute content, they apply sophisticated techniques to filter out influential content and people. Facebook’s EdgeRank algorithm (Seaver, 2019) and Twitter’s Timeline ranking (Koumchatzky and Andryeyev, 2017) decide what users will see on their timelines/feeds. The popularity of a person or post (measured in terms of number of likes, followers or

<sup>7</sup> These decisions are not completely autonomous since they are influenced by platform decisions and algorithms as we will discuss later.

retweets) is indirectly included in the model's prediction of how interesting and engaging a Tweet would be specifically to you. Third, connectivity is the ability of social platforms to connect content to users and advertisers. Again, connectivity has a human and an algorithmic component. People are either directly (e.g. friends) or indirectly (e.g. liking the same product) connected to other people and these connections form the basis of automated recommendations and personalization. Lastly, datafication is the ability to infer valuable information based on customers' demographic or profiling data, likes, shares, friends, and so on. Initially, data was considered as a by-product of social networks, but as the platforms matured, data and predictive analytics have become the main asset in their monetization model.

The elements of the social media logic shape human interactions and information flow on social networks and are crucial in understanding how social media affects political communication, public opinion forming, and the well-functioning of our democracy.

### 1.2.2 *From liberation to repression*

To understand how this new model of information flow can become both an opportunity and a challenge for democracy we need to understand the effects of the two vital changes we just discussed. The first element enables citizens to find and distribute views that are normally excluded from political discussions in the mainstream media. On social media, one can easily seek groups of like-minded people, coordinate collective action, and support political candidates. This allows citizens to team up for a joint cause, hold governments accountable or fight for minority rights (Bennett and Segerberg, 2012). The Arab Spring is a well-known example of how social media helped the people to protest oppressive regimes in the Middle East and North Africa (Howard et al., 2011). More recently, social media played an important role in the Black Lives Matter protests all over the world (NBC News, 2020). Yet, social media can amplify extreme and "anti-democratic" voices as well. It is easier for those with minority views to find like-minded people, and at the same time, the fact-checking roles of journalists and media diminish, clearing the path for controversial ideas to spread. The very same mechanisms that allow citizens to fight for wider political inclusion, are now empowering terrorist groups (Gates and Podder, 2015), far-right extremists (Daniels, 2018), and conspiracy theorists (Amarasingam and Argentino, 2020) that challenge core democratic values.

The second element allows politicians and political parties to shape their content and communicate directly with voters. This turns social media into a powerful tool in the hands of populist parties and candidates, as they can get large exposure through provocative statements. In Belgium (Flanders) for example, Tom Van Grieken has put his political party Vlaams Belang successfully back on the trails with a well thought-out social media strategy. Their social media posts focus on engaging content and contain hyperbolic claims that spark outrage. Often they requested their

audience explicitly to share the posts so that “the truth could not be covered up” (Maly, 2020). More dangerously, authoritarian regimes can exploit the information freedom to silence others. Censorship and propaganda tactics are applied to amplify the regime’s messages while silencing the opposition’s. Next to repression strategies to undermine freedom of information online (such as intimidation and jailing but also internet shutdowns, removal of online content, and algorithmic manipulations) authoritarian regimes have harnessed the openness of social media platforms by “flooding” platforms with propaganda. They can use paid trolls or automated bots to promote government propaganda or flood antiregime protest hashtags. They may even spread misinformation, or harass regime opponents online (Tucker et al., 2017).

In short, social media can be seen as both a medium of liberation as well as a technology of repression. Then how should this technology be regulated to prevent anti-democratic actions without restricting the democratic nature of the platforms themselves? To answer this question we need scholars, policy makers, civil society, journalists, and political actors to reflect on the responsibilities of government, platforms and citizens in the new digital age (Persily, 2017). Scholars from various domains can add to this debate by unraveling the complex interactions between social media and politics. Doing this requires innovative ways to study the massive amounts of digital and social media data that are available.

### 1.3 Computational social science

A day in the life: We wake up in the morning and check our email or scroll through Facebook. We track our morning run with Strava, read the news online while having breakfast, consult the weather forecast on our mobile device, and message our friends via Whatsapp to meet for coffee. We bike there using Google maps and we pay for our coffee by credit card. Afterwards, we rate the coffee bar on Tripadvisor and post a picture on Instagram, which we then need to open every five minutes to keep track of the number of likes we have gained. Before noon we’ve left behind a whole path of digital breadcrumbs which, when pulled together, offer an increasingly comprehensive picture of individuals and groups. This creates the potential of manipulating us for commercial or political purposes, but also to expand our understanding of human behavior and society in a way that was simply unimaginable just a decade ago (Lazer et al., 2009).

#### 1.3.1 *Two research paths*

Social science is the scientific study of human relationships and society. Computer science is the study of computers and computational algorithms for processing information (Belford and Tucker, 2020). At the intersection of social science and computer science lies Computational Social Science (CSS), or “the study of social

phenomena using digitized information and computational and statistical methods” (Wallach, 2018, p. 1). Both fields differ fundamentally with respect to three points: goals, methods and data (Wallach, 2018).

**GOAL** A fundamental reason for many misunderstandings between social scientists and computer scientists is their primary research goal. An important research area in the field of computer science is concerned with the task of prediction: using observed data for the purpose of predicting unseen or future observations (Shmueli et al., 2010). Machine learning has traditionally focused on prediction tasks such as classifying images, natural language processing, recognizing handwriting, targeted advertising, and predicting fraud. Social scientists on the other hand focus on explanation. The goal is to understand why and how we observe certain social phenomena. Social scientists are trained to test hypotheses from observed data, given a theoretical model (Shmueli et al., 2010).

To illustrate the tension between these two goals, consider the following example. In text analysis, computer scientists typically apply methods for automated content analysis that are optimized to classify individual documents. For example, a machine learning model can automatically detect whether a text is about social welfare, environment or macroeconomics. Social scientists instead rather want generalizations about the collection of documents. For example, what is the proportion of texts that is about the environment. A model that is optimized for accurate individual classification does not necessarily perform well to estimate category proportions. Suppose that all errors happen in the same category (e.g. some of the texts about the environment are incorrectly classified as macroeconomics), then the statistical bias in estimating the aggregate proportions could be very high. This is often no direct problem from a prediction perspective, since individual document classification performs well, but will lead to biased results when interested in aggregated category proportions. On the other hand, a method optimized for estimating document category proportions can provide unbiased estimates of category proportions even when the individual classifier performs poorly (Hopkins and King, 2010).

**DATA** Computer scientists typically work with large-scale (digitized) observational or behavioral data, since these massive datasets are extremely useful for accurate prediction (De Cnudde, 2017). In contrast, social scientists are used to work with data carefully collected to answer specific questions, such as survey or experiment results. Within the context of this thesis, we will refer to the first type of data as *digital trace data*,<sup>8</sup> the records of activity we leave behind in the digital world. Digital traces are a form of *fine-grained behavioral data*, typically large datasets that document individual behavior at a fine-grained level such as payment data, location data or website visits. These data have proven valuable in predictive applications such as financial credit scoring (De Cnudde et al., 2019; Tobback and Martens, 2019), fraud

<sup>8</sup> We will refer to this data as behavioral data and digital trace data interchangeably

detection (Vanhoeveveld et al., 2019, 2020), and churn prediction (Verbeke et al., 2014). This type of data is sparse and high-dimensional: the total number of possible actions (e.g. pages to like) is huge but because of limited “behavioral capital” any individual can only take a very small fraction of all the possible actions (Junqué de Fortuny et al., 2013; Tibshirani et al., 2015). This implies that every individual or instance will have zero values for many of the features. Although textual data show some differences from behavioral data (De Cnudde, 2017), they are often considered similar, since they are both sparse and high-dimensional.

On the other hand, for what we will call *traditional data*, it is the other way around: the data is dense, every instance receives a value for every feature and the number of features is generally lower than the number of samples or respondents (Junqué de Fortuny et al., 2013; Tibshirani et al., 2015). This is the case for traditional survey data, where each individual will have a value for almost every question. An example of both data types is given in Figure 1.3. While survey research has been at the heart of social science for decades, social scientific research with digital trace data, including social media data, has been growing rapidly in the last few years (Stier et al., 2019). In this thesis we will analyze social interactions and textual data on Twitter, and public page likes on Facebook.

	Gender	Age	...	Education
Person 1	M	25	...	high
Person 2	F	55		low
...				
Person n	F	30		high

(a) Traditional (survey) data

	Radiohead	Ben & Jerry's	...	Greenpeace	...	Fightclub
Person 1	1	0	...	0	...	0
Person 2	0	1		1		0
...						
Person n	1	0		0		1

(b) Behavioral (Facebook) data

Figure 1.3: An example of traditional data (a) and fine-grained behavioral data (b). For the behavioral (Facebook Like) data the value will be 1 if the user liked the page and 0 if not.

When collecting data, the sample is required to be representative for the population under study in the case of explanatory modeling, since the goal is to draw statistical inference about the parameters for the population based on this sample (Nagler and Tucker, 2015). For prediction, the data used to train the prediction model needs to be representative for the unseen data (Schat et al., 2020). Furthermore, the size of the sample is crucial to increase prediction performance. Especially when using fine-grained behavioral data, we continue to see improvements in predictive performance when including more data (Junqué de Fortuny et al., 2013).

**METHODS** Since the goal of prediction and explanation is different, they lead to different modeling approaches. Explanatory modeling requires interpretable statistical models, that are easily linked to the underlying theoretical model (Shmueli et al., 2010). Often, regression-type methods are used to test hypotheses, where the coefficients of the model are the objects of interest (Cranmer and Desmarais, 2017). Interpretability is not a strict requirement for good prediction, although it is often desired and has received increasing research attention (Ramon et al., 2021; Burkart and Huber, 2020). Next to statistical models, complex data mining algorithms such as neural networks can be applied for accurate prediction, without shedding light on the underlying explanatory mechanisms. Some methods exist that are very useful for prediction, but not for explanation. In Chapter 3 and 5, we will make use of regularization: an extra constraint introduced to the optimization function that penalizes the weights of coefficients. This technique has the effect of reducing variance at the expense of introducing bias to the model. It is especially useful when modeling sparse data: simply optimizing the standard objective function would lead to estimates that will overfit the data (Tibshirani et al., 2015).

A final important distinction is model evaluation. Validating a explanatory model involves goodness-of-fit tests (e.g., normality tests) and model diagnostics such as residual analysis. In predictive modeling, the biggest danger to generalization is overfitting the training data. Hence, prediction performance is evaluated on a part of the data that was not used for training the model (a test set). I will introduce some frequently-used prediction methods and evaluation metrics in the next chapter.

**THE CROSSROADS** The fundamental goal of social science (explanation) and computer science (prediction) is different, which will lead to practical implications to each step of the modeling process, including the way data is selected, the type of models that are used, and model evaluation.<sup>9</sup> Understanding these differences is essential to develop a rigorous research design. However, rather than opposites, explanation and prediction can be considered as two dimensions so that every model will possess some level of each (Shmueli et al., 2010). A well-specified explanatory model is expected to make accurate predictions on unseen instances and should have some level of predictive power. Similarly, it can be argued that a well-performing predictive model should be congruent with theory and demonstrate some explanatory power as well. In this thesis, I will apply predictive modeling to contribute to explanatory understanding. Predictive models can contribute to theory by uncovering patterns in large sets of data that might not have been obvious to the theory-building analyst (Cranmer and Desmarais, 2017). I will apply the term *predictive* relationship, to refer to a data-driven correlation discovered using a predictive modeling approach. In contrast to an *explanatory* relationship, which is derived from theory and confirmed by the data. Although ideally both types of relationships are equivalent, this is not guaranteed due to practical differences in the way they have been obtained.

<sup>9</sup> As discussed above. For more detail, I refer to the paper of Shmueli et al. (2010)



### 1.3.2 *Opportunities for empirical research*

Clearly, the application of computer science methods to study massive datasets of human behavior offer new opportunities to answer social science questions. With traditional surveys, specific questions can be asked about the phenomena of interest in a structured way. This data is analyzed for a sample of the population to draw statistical inference about the parameters for the population (Nagler and Tucker, 2015). Behavioral data on the other hand, can be seen as completely open-ended responses for which there was not even asked a question. Therefore, they can provide an unfiltered look and unique information that is difficult to grasp using survey techniques. Moreover, the data can be collected and mined at a scale previously unknown to social scientists. More specifically, the use of digital trace data has three important advantages over the use of more traditional measurement approaches (Jungherr and Theocharis, 2017).

First and most prevalent, self-reported behavior has the potential to be biased due to ill-phrased questions, social desirability or the inability of respondents to recall facts correctly (Furnham, 1986). Several studies have discovered substantial discrepancies between objective and self-reported behavior, due to recall error or subjective interpretation of the questions (Prior, 2009; Guess et al., 2019a; Mosleh et al., 2020). Additionally, behavior observed in designed experiments is difficult to generalize to the real world due to the artificial conditions the subjects are put in. Digital trace data emerge from user behavior in a more natural setting.

Second, collecting survey data is time-consuming and costly. Filling out a survey can be a large effort for the respondents. For this reason, researchers struggle to find a representative number of participants putting survey research under pressure. Behavioral data on the other hand, measures a certain behavior directly and requires little effort from the participants. The sheer size digital trace data are supplied, is unprecedented in the social science field. Possible ways to collect digital data are (free and paid) Application Program Interfaces (APIs) provided by digital platforms, web scraping, or direct research collaborations with the data sources (Freelon, 2018; Perriam et al., 2020). The relative ease of collection and size of the data allow researchers to compare patterns between subgroups with small relative weights, or to identify small effect sizes that would otherwise be indistinguishable from noise (Jungherr and Theocharis, 2017).

Lastly, digital traces offer a very detailed and fine-grained look on user behavior and interactions, at a resolution that would be difficult to attain with traditional surveys (Monroe et al., 2015). Data mining algorithms allow researchers to discover previously unknown patterns in these massive amounts of data which can help them to reach a better understanding of social processes (Lazer et al., 2009).



### 1.3.3 Challenges

Despite the enormous opportunities of digital trace data for empirical research, translating the “measurement revolution” into valuable social science knowledge remains a technical and conceptual challenge (Jungherr and Theocharis, 2017).

Importantly, behavioral data is only a proxy for the actual behavior of individual citizens, it only measures what happens on the service or device that measured the data. For example, what you like on Facebook is not necessarily what you like in real-life and thus true behavior can only be studied indirectly (Dalton, 2016). This might not be a problem if you are interested in identifying usage patterns on selected platforms, but drawing inference on larger social phenomena and generalizing results from this type of data must be done with great caution (Nagler and Tucker, 2015; Dalton, 2016). For instance, several studies have analyzed selective exposure on social platforms (Bakshy et al., 2015; Barberá et al., 2015; Guess et al., 2018; Liu et al., 2016), without necessarily intending to generalize results to the offline world. Additionally, digital traces are influenced by the technological features of platforms they come from, such as their rules-of-conduct and algorithms. In Section 1.2.1 we discussed the social media logic, the processes through which platforms channel social traffic. Neglecting these mediating factors can lead to wrong conclusions about social phenomena. Digital trace data do not necessarily measure the behavior of interest, which has been called the “mirror fallacy” (Jungherr, 2018).

Second, even though indeed massive data collections exist, access to digital trace data by researchers depends on the provisions of the corporations holding the data. Some social media platforms have a free and convenient-to-use Application Program Interface (API) that allow researchers to collect data (e.g. Twitter<sup>10</sup>). However, the closure of Facebook’s API for researchers in 2018 points out the importance of having alternative ways to collect data such as web scraping (Freelon, 2018) or direct research collaboration with data sources, while keeping in mind the legal and ethical limitations. For all above three methods however, researchers only gain access to a selection of the data, determined by the access policies of the data holders. How the collected data is extracted from the underlying complete dataset is often unknown (Jungherr, 2018).

Third, misrepresentation bias —the sample that is studied does not accurately represent the population— can occur with traditional data through self-selection of the respondents, erroneous selection decisions by the researchers, or high non-response rates (Heckman, 1979). Digital trace data is inherently biased because the digital population is not representative for the general population (Mellon and Prosser, 2017).<sup>11</sup> As we described in Section 1.1, the user-base on social media consists mostly of young, male and highly educated people. Yet, the exact composition seems

<sup>10</sup> <https://developer.twitter.com/en>

<sup>11</sup> Of course, this is only a problem if our goal is to study the general population. If our goal is to study behavior on a social platform, the sample only needs to be representative for the platform’s audience.

Table 1.1: Social science versus computer science.

	Social science	Computer science
Goal	Theory-driven explanation of social phenomena	Individual prediction
Examples	Survey data to study voting behavior (Berelson et al., 1954; Campbell et al., 1960)	Predict political preference from Facebook data (Kosinski et al., 2013; Kristensen et al., 2017)
Data	Traditional data (surveys and experiments)	Behavioral data (digital trace data)
Models	Interpretable statistical models	Data mining algorithms
Challenges	1) Response bias 2) Time and effort 3) Misrepresentation bias	1) Mirror fallacy 2) Data access 3) Misrepresentation bias 4) Privacy and security

to fluctuate over time, making it hard to use weighting schemes to correct for the skewness (Schober et al., 2016). Again, this is not a problem if you are interested in the platform audience, but complicates generalization to the offline population.

A last issue with behavioral data is privacy. Since this data is often very personal and sensitive, it must be collected and stored with respect for users' privacy, addressing challenges on user consent, data anonymization, secure storage, etc. (Zimmer, 2010). The next section will be devoted to an extensive discussion of the privacy and ethical concerns of conducting research with online, human-generated data. Of course, these ethical challenges are not entirely new, but the rapid increase in available online data has raised new and unexpected ethical questions.

An overview of the key differences between social and computational science and their respective challenges to study human behavior and politics can be found in Table 1.1. It should be clear that generating valuable insights from digital trace data is challenging. Simply throwing data mining algorithms at big data does not lead to identification of new laws of social life. Yet, because of the enormous amount of information available in behavioral data, there is much to learn from it when following a rigorous research approach (Nagler and Tucker, 2015; Dalton, 2016).

#### 1.3.4 *Privacy and ethical concerns*

Facebook and other online environments provide tremendous opportunities for social science research to study human behavior. At the same time, the rapid increase in available online data and technological progress raises new and unexpected ethical questions related to online collection, storage, and use of human subjects' data (Zook et al., 2017). These questions become urgent as the data and research possibilities move well beyond those typical in the social sciences, to more directly address sensitive aspects of human behavior and our daily lives.

Even seemingly benign data can contain sensitive or private information and has the potential to impact people's lives (Zook et al., 2017). Plenty of examples show how publicly available Twitter data can reveal things about you that you might not usually want to share with strangers, such as your personality (Golbeck et al., 2011), political orientation (Cohen, 2013; Barberá, 2015; Liu et al., 2016), emotions (Colnerić and Demsar, 2018) and mental health (Coppersmith et al., 2014).

Privacy is the right to control information about yourself and the way it is communicated to others (Westin, 1968). It is a universal human need and right (UN General Assembly, 1948; Council of Europe, 1950), and has taken a prominent role in the digital age. Privacy invasions have been related to harmful activities such as surveillance, disclosure, exposure, intrusion, appropriation, blackmail, and decisional interference, thus damaging an individual's dignity, or even freedom (Solove, 2005).

From a legal standpoint, the introduction of the European Union's General Data Protection Regulation (GDPR) in May 2018, offers a regulatory framework on data protection and privacy for researchers collecting (online) personal data. GDPR defines personal data as "any information which are related to an identified or identifiable natural person" (Art. 4(1) European Parliament and Council, 2016). GDPR allows to process such data under certain conditions, among which —most relevant for research— unambiguous consent by the data subject for one or more specific purposes, or, performance of a task carried out in the public interest (Art. 6 European Parliament and Council, 2016). Article 5 of the regulation provides principles of personal data processing relating to (a) lawfulness, fairness, and transparency; (b) purpose limitation; (c) data minimisation; (d) accuracy; (e) storage limitation; and (f) integrity and confidentiality. As a general rule, researchers processing (online) personal data should be abreast of these regulations.

Next to regulation, there are also other terms and conditions researchers need to comply with, such as the platform policy or specific requirements from their research institutions or funding bodies. Finally, after compliance with all relevant terms and conditions an ethical reflection of the researchers is needed. Various instances are updating ethical guidelines for Internet research and human subjects (including Franzke et al., 2020; Office for Human Research Protections and Services, 2020).

Yet, the variety in data sources, research topics, and methodological approaches complicates the draft of universally applicable ethical guidelines. Rather, all researchers engaging with (data from) human subjects, have the ethical responsibility to minimize potential harm. Next, we will discuss four key areas of concern that are specific to internet and social media data (Townsend and Wallace, 2016).

**PRIVATE OR PUBLIC** All activity on Twitter is public.<sup>12</sup> On Facebook, users can adapt their privacy settings and choose to share content either publicly or with friends only.<sup>13</sup> Information that is not shared publicly is clearly private, and researchers need explicit user consent to access that data. Publicly available information can technically be obtained by researchers without asking for consent. However, simply because the data is publicly available it does not free us from ethical considerations. Intuitively, it would not be reasonable to assume that users would agree to have their social profile broadcasted on national television as a high risk profile for mental health issues.<sup>14</sup>

Whenever collecting and storing personal information—even if it’s publicly available—you are bound to ethical and even legal requirements (such as GDPR in Europe and the data providers’ policies) (Martens, 2021). Twitter provides researchers access to (parts of) their data through an API, but the use of this services requires compliance with their developers policy which restricts certain use cases.<sup>15</sup>

**INFORMED CONSENT** Informed consent implies that the person providing consent understands what data will be used and for what purpose (Barocas and Nissenbaum, 2014). In more traditional social science research this is usually built into an informed consent form that participants need to agree with before participating in an experiment or survey. For the private Facebook data (Public Page Likes) we have collected as part of this PhD research, we followed this approach (see Chapter 2). Participants had to agree to our privacy policy, in which we clearly explained which data will be collected, how it will be stored, and what we will do with it—we explicitly mentioned predicting political preference since this is sensitive information. Importantly, one should only use data from people that consented. It is not allowed to collect data from the broader network. This is less obvious than it seems, because this also applies to posts that are posted on a timeline. The user might agree to share their timeline data with researchers, but not everyone who posted on the timeline might agree.

<sup>12</sup> <https://twitter.com/en/privacy>

<sup>13</sup> [https://www.facebook.com/help/120939471321735?helpref=faq\\_content](https://www.facebook.com/help/120939471321735?helpref=faq_content)

<sup>14</sup> In fact, it is prohibited by Twitter developers terms to derive or infer health information about a Twitter user (<https://developer.twitter.com/en/developer-terms/more-on-restricted-use-cases>)

<sup>15</sup> <https://developer.twitter.com/en/developer-terms/agreement-and-policy>

Acquiring informed consent becomes even more problematic for large public datasets, such as Twitter data. On Twitter, users have implicitly agreed to the access and use of their data by third parties, including researchers, by accepting the platform's policies.<sup>16</sup> Yet, this policy is seldom read in detail (Bakos et al., 2014), and is not specific about the research purpose. Therefore, users cannot unambiguously be considered to have given their informed consent. Kosinski et al. (2015) propose the following conditions that justify the use of public data without seeking explicit consent: (1) It is reasonable to assume that the data were knowingly made public by the individuals; (2) Data are anonymized after collection and no attempts are made to de-anonymize them; (3) There is no interaction or communication with the individuals in the sample; and (4) No information that can be attributed to a single individual, including demographic profiles and samples of text or other content, is to be published or used to illustrate the results of the study. We could add to this that researchers should be extra cautious of potential risks when they aspire to infer sensitive information from public data, such as political affiliation, financial status, ethnicity, health, etc.

One possible concern with the requirement to ask explicit consent from data subjects is that it will bias the sample, considering privacy-savvy users are less likely to consent. Under the GDPR, there are different privacy standards for EU and non-EU citizens, leading to different research possibilities for both groups (Greene et al., 2019).

**ANONYMITY** Personal data should be stored without personally identifiable information (e.g., name, email address, IP address) to protect user privacy. This does not guarantee anonymity, because reidentification is still possible. Behavioral data can disclose individuals in unanticipated ways. For example, in 2006 Netflix revealed a massive dataset of around 100 million movie ratings, intended for a data mining contest to improve their movie recommendation algorithm. Although the movie ratings Netflix published were anonymous, researchers could identify several Netflix users by comparing the Netflix data to the Internet Movie Database (IMDb). Movie preferences are sensitive information since they can reveal an individual's personal interests, including sexual orientation, mental illness, etc. Netflix was forced to cancel the contest after a lawsuit accused them of exposing movie preferences of their users (Singel, 2009). Plenty of other examples show how "anonymous" publicly available internet data can lead to unintentional data leakages (Barbaro et al., 2006; Zimmer, 2010; Schwartz et al., 2017). Privacy-preserving techniques, such as differential privacy or k-anonymity, can reduce the likelihood of reidentification (Martens, 2021). Yet, making deidentified sensitive information publicly available always involves a risk and might better be avoided. This way, personal data privacy standards —albeit

<sup>16</sup> Twitter's privacy policy states "To facilitate the fast global dissemination of Tweets to people around the world, we use technology like application programming interfaces (APIs) and embeds to make that information available to websites, apps, and others for their use - for example, displaying Tweets on a news website or analyzing what people say on Twitter." (<https://twitter.com/en/privacy>)

absolutely necessary for individual protection—complicate scientific reproducibility of research results.

**RISK OF HARM** Risk of harm is greater when dealing with sensitive data, which when revealed might lead to embarrassment, reputational damage or even discrimination. A researcher’s responsibility towards their participants and the protection of their data increases when the risk of harm or vulnerability of the participants increases (Townsend and Wallace, 2016). Furthermore, even studies with good intentions, such as using educational data to identify potential dropouts, can lead to unintentional harm by labeling students as “failures” (similar to crime analytics) (Shmueli, 2017). Finally, there is always an indirect risk of technological and theoretical advancements to be applied in harmful ways. Knowledge developed by social and computer scientists to better understand and predict human behavior can be abused to manipulate individual’s beliefs and behavior. Big data analytics and microtargeting have the potential to control voter behavior and public opinion forming (Gorton, 2016).

#### 1.4 Social media to study political behavior

We will illustrate the opportunities and challenges of computational social science and social media data with some examples in the field of political science. We’ll first discuss social media data to study behavior of political actors, followed by behavioral studies of the larger public (cfr. Politics and Citizens in Figure 1.2a). Table 1.2 provides an overview of the research applications discussed below.

**POLITICAL ACTORS** Traditionally, political research would examine party manifestos, campaign ads or press releases to study strategic communication choices by political actors (Tresch et al., 2017). However, nowadays social media represents an interesting alternative, as it is perhaps the most widely accessible form of party communication, with higher temporal adaptability and interaction potential (De Sio and Lachat, 2020). There is growing scholarly interest in parties’ communication on Facebook and Twitter providing insights in issue attention (Vargo et al., 2014; Van Dalen et al., 2015; Peeters et al., 2019; Van Ditmars et al., 2020), campaign strategies (Nulty et al., 2016) and populism (Stier et al., 2017). Especially Twitter is increasingly used by political parties and politicians to communicate with citizens, opinion leaders and journalists (Jungherr, 2016; Vargo et al., 2014).

Besides content, also the structure of interactions between politicians is of interest. Peer influence in social networks is known to affect opinions and attitudes. Again, social media provides a rich data source for studying interpersonal relations, e.g. followers on Twitter and friends on Facebook. Several studies have analyzed parliamentary social networks to gain insights in political polarization (Conover et al.,

2011; Esteve Del Valle and Borge Bravo, 2018), opinion leadership (Borge Bravo and Esteve Del Valle, 2017), the underlying structure of political groups and countries (Cherepnalkoski and Mozetič, 2016), etc.

Aside from social media as just another tool for politicians to communicate their message and position themselves, social media can also be used as an anti-democratic tool by less-democratic regimes, as discussed in Section 1.2.2. In that case, social media data can help to better understand digital repression strategies by authoritarian actors. Research examples include online censorship (King et al., 2014), regime response to protests (Munger et al., 2014) and bots and troll detection (Stukal et al., 2017).

**CITIZENS** Because of the advantages of social media data, discussed in Section 1.3.2, it offers new opportunities to measure voter behavior. Some of the earliest work in this field focused on measuring public opinion about political issues. Several studies have attempted to predict aggregate election results and overall political sentiment (Barclay et al., 2015; Giglietto, 2012; Ceron et al., 2015; MacWilliams, 2015). Social media is a valuable complement to survey data, as an early indicator of changes in public opinion and insights on unpolled topics (Barberá and Steinert-Threlkeld, 2020). Research is also concerned with the effects of social media on different forms of online and offline political participation. Scientists attempt to understand how information and social interactions on social media influence voter turnout (Settle et al., 2016; Jones et al., 2017; Bond et al., 2012). Next to that, social media are valuable in organizing and mobilizing collective action and protests (Bennett and Segerberg, 2012; Enikolopov et al., 2020).

Other than aggregate measurements and prediction, Kosinski et al. (2013) showed that Facebook Like data can be used for predicting individual personality traits and political attitudes. Since then, other researchers affirmed the potential of social media data for predicting individual political orientation (Bond and Messing, 2015; Barberá, 2015; David et al., 2016; Kristensen et al., 2017; Chiu and Hsu, 2018; Bach et al., 2019).

Another well researched behavior on social media is homophily and polarization. Decades of sociological research have demonstrated that individuals who are closely connected in a social network tend to show similar behaviors and characteristics. This is explained by the self-reinforcing dynamics of two well-documented social processes: homophily, or the tendency of individuals to associate and bond with similar others (McPherson et al., 2001), and social influence, the process in which individual behaviors are influenced by exposure to others (Kelman, 1958). *Political* homophily refers to the tendency to associate with others who are similar in political ideology. Greater political homophily implies decreased chances of politically diverse interactions and increased rates of ideologically similar ones, which tends to reinforce personal political views and in-group commitment (Boutyline and Willer, 2017). This



might potentially lead to polarization. Several authors have studied the degree of polarization among the electorate (Bond and Messing, 2015; Moeller et al., 2018), the existence of echo chambers (Eady et al., 2019; Barberá et al., 2015), exposure to (ideologically diverse) news and opinions (Wells and Thorson, 2017; Bakshy et al., 2015; Bail et al., 2018; Messing and Westwood, 2014), and the relation of these mechanisms to the spread of misinformation and fake news (Guess et al., 2019b; Allcott et al., 2019), and harmful speech (Siegel et al., 2019; Siegel and Badaan, 2020; Müller and Schwarz, 2020).

### 1.5 Contribution and framework

Through five different studies, we contribute theoretically and methodologically to the existing literature on social media and politics. Though the theoretical contributions lie in different subfields (see table 1.2), our exploratory approach, based on the Computational Social Science methodology discussed in Section 1.3, is what unites these works. Research starts with a theoretical question grounded in social or political science. The question is either related to political behavior on social platforms or behavior that can be measured through social media data. In the next step, behavioral data is collected from the relevant platforms and potentially combined with survey data. Data mining algorithms are applied to the data to discover (previously unknown) patterns. This step often leads to incremental methodological improvements. Finally, new insights from the data can help to reach a better understanding of the theoretical processes and lead to theory building. All this is done with respect for privacy and ethical considerations (see Figure 1.4).




Figure 1.4: Exploratory research methodology

In **Part I** of this thesis we study the **behavior of political actors on Twitter**. Twitter allows politicians and political parties to communicate directly and autonomously with voters and the media. It is therefore not evident that established theories in political communication also apply to this medium. In **Chapter 3** we contribute to the rapid increase of studies that deal with different aspects of party communication. We propose an exploratory approach to analyze issue communication that is adapted



Table 1.2: Examples of research applications in social media and politics

Actor	Category	Examples
Politics <i>PART I</i>	Communication <i>Chapter 3</i>	Issue attention (Vargo et al., 2014; Van Dalen et al., 2015; Peeters et al., 2019; Van Ditmars et al., 2020), campaign strategies (Nulty et al., 2016), and populism (Stier et al., 2017).
	Social interactions <i>Chapter 4</i>	Political polarization (Conover et al., 2011; Esteve Del Valle and Borge Bravo, 2018), opinion leadership (Borge Bravo and Esteve Del Valle, 2017), the underlying structure of political groups and countries (Cherepnalkoski and Mozetič, 2016)
	Repression	Online censorship (King et al., 2014), regime response to protests (Munger et al., 2014) and bots and troll detection (Stukal et al., 2017)
Citizens <i>PART II</i>	Political affiliation <i>Chapter 5</i>	Predicting individual political orientation (Bond and Messing, 2015; Barberá, 2015; David et al., 2016; Kristensen et al., 2017; Chiu and Hsu, 2018; Bach et al., 2019).
	Homophily and polarization <i>Chapter 6 and 7</i>	Polarization and echo chambers (Bond and Messing, 2015; Moeller et al., 2018; Eady et al., 2019; Barberá et al., 2015), exposure to opposing views (Wells and Thorson, 2017; Bakshy et al., 2015; Bail et al., 2018; Messing and Westwood, 2014), misinformation and fake news (Guess et al., 2019b; Allcott et al., 2019), harmful speech (Siegel et al., 2019; Siegel and Badaan, 2020; Müller and Schwarz, 2020)
	Public opinion and political participation	Election results and overall political sentiment (Barclay et al., 2015; Giglietto, 2012; Ceron et al., 2015; MacWilliams, 2015), voter turnout (Settle et al., 2016; Jones et al., 2017; Bond et al., 2012), protests (Bennett and Segerberg, 2012; Enikolopov et al., 2020)

to the volatility of social media text. This method is applied to six Belgian (Flemish) political parties on Twitter and we show that it helps unravel how political parties profile themselves on Twitter and which strategies are at play. This in turn contributes to classical literature on issue competition, party unity, and party discipline.

In **Chapter 4** we study political communication from a network perspective. The aim of this research is to better understand how social media affect the communication flow among political actors (parliamentarians) and what the influence of institutional context is on this network behavior. We collect one year of Twitter data from all members of parliament and government in 12 countries. Social network analysis (SNA) is applied to analyze the relation between network properties of the parliamentary Twitter networks and the political system and democratic functioning of the countries. Secondly, we analyze the inter-party communication and its link to party ideology. We demonstrate the importance of institutional context and the opportunities of social network analysis on Twitter for comparative research.

Secondly, in **Part II**, the **behavior of the public** is the object of research. Evidence of growing political polarization, especially in the United States, invites speculation about whether political polarization extends to every aspect of our daily life. However, detailed information on individuals' lifestyles is very difficult to collect, which complicates empirical and comparative studies in this domain. In **Chapter 5** we explore the potential of Facebook Likes to complement traditional survey data and study the interrelation between ideology and lifestyle choices. We collect a unique dataset of Facebook Likes and survey data of more than 6,500 participants in Belgium, and infer the political and ideological preference of our respondents. The results indicate that non-political Facebook Likes are indicative of political preference and are useful to describe voters in terms of common interests, cultural preferences, and lifestyle features.

Subsequently, we utilize Facebook Like data to test whether polarization permeates society or if it is more limited to strictly political domains and among politically active individuals. In **Chapter 6**, we combine survey and Facebook Like data from more than 1,200 respondents in the United States. Our evidence adds nuance to the narrative of widespread polarization across lifestyle sectors, and it suggests domains in which cross-cutting preferences are still observed in American life. Finally, **Chapter 7** compares polarization in the two-party American system to the multi-party Belgian case. We find ideological divides in political domains to be much less outspoken in Belgium. Figure 1.5 summarizes the contributions of this PhD thesis with respect to political behavior in the social media age.

In conclusion, **Chapter 8** summarizes this work and presents avenues for further research. We suggest recommendations to enable more impactful research in the field of computational social science. Finally, we end with an outlook on the future role of social media in our society.

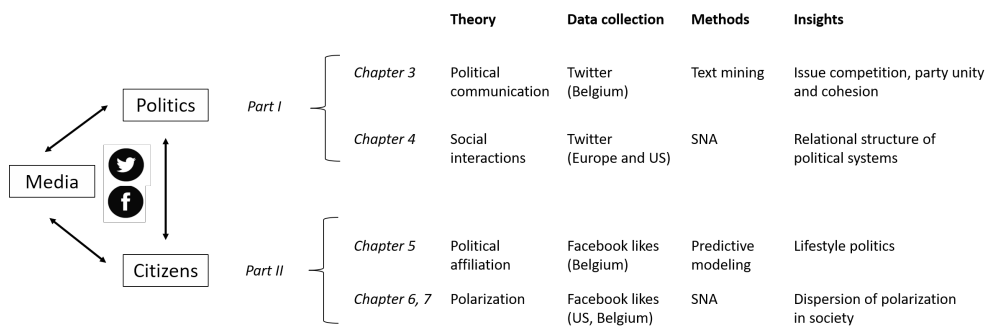


Figure 1.5: Contributions of this PhD thesis



## Data and methods

---

This chapter presents the data and methods used throughout this thesis, and is intended to increase the readability of the following chapters. The technical details of the data collection from the Twitter and Facebook platform will be explained. Next, I will examine the ethical implications of our data collection and research, based on a framework that summarizes the ethical considerations that were discussed in the introduction. Finally, a short introduction to relevant data mining concepts and performance metrics is provided; and the techniques that will be used in the following studies are presented. All code used throughout this thesis is written in Python or R, and is made available on GitHub.<sup>1</sup>

---

<sup>1</sup> <https://github.com/SPraet>

## 2.1 Data collection

The first part of this thesis focuses on Twitter data, the second on Facebook data. Both were collected through an Application Programming Interface (API) provided by the platform. An API is an interface through which third-party developers can connect with the platform. The API can be used by researchers to collect data from the social media platform for empirical analysis. As mentioned before, Twitter data is publicly available while the data we collected from Facebook (Public Page Likes) is private. Therefore, the data collection process for both datasets is different in terms of the API application process, user consent, and the type and amount of data that can be collected, as are the ethical considerations. We base an ethical discussion of our research and data collection on the data science framework of Greene et al. (2019). The framework unites GDPR regulation and moral concerns, and consists of the following steps: collecting data, using data, sharing data, generalizability, and communication. The following sections discuss how the Twitter and Facebook data for this thesis were collected, and which ethical considerations were involved.

### 2.1.1 Twitter data (Part I)

There are APIs available on the Twitter platform that software developers can engage with to search or stream tweets, send direct messages, embed Twitter within their website, create and manage ads, etc. The Standard search API<sup>2</sup> allows to search a sampling of recent tweets published in the past seven days.<sup>3</sup> To collect tweets over a longer period of time, it is recommended to stream realtime Tweets, using the Streaming API.<sup>4</sup> The only steps that are required before you can start data collection through one of the provided APIs is to create a Twitter account and apply for a developer account, which is usually approved within less than two weeks.<sup>5</sup> A list of politicians and their Twitter accounts was manually created to start streaming the tweets of the relevant accounts for the desired time period. I used the python Twitter API wrapper Tweepy to stream tweets.<sup>6</sup>

Twitter data is publicly available, this simplifies the collection process, but does not allow us to do whatever we want with the data. We are still bound to the Twitter policy on data collection and use, and ethical limitations. We only collect Twitter data from politicians or parties. Public figures can reasonably be expected to know their Twitter information is public, as they consciously use Twitter to communicate and engage with the public. If a politician has two accounts we only include the one that is used for their political function. No politician will be surprised that their communication is studied in academic research, and the purpose of our studies is

<sup>2</sup> <https://developer.twitter.com/en/docs/tweets/search/overview>

<sup>3</sup> Just recently, Twitter announced they will open up their full tweet archive to academic researchers in the new free Academic Research product track (Tornes and Trujillo, 2021)

<sup>4</sup> <https://developer.twitter.com/en/docs/tweets/filter-realtime/overview>

<sup>5</sup> or at least this was the case at the time I created a developer account, by the end of 2017.

<sup>6</sup> <https://docs.tweepy.org/en/v3.4.0/streaming.html>

not beyond what can reasonably be expected from academic research. Results are only reported at the party level and we never include individual examples. For these reasons, we do not require explicit user consent to collect personal data. In accordance with the Twitter policy, we only share tweet IDs. Based on the tweet ID, the full tweet can be collected by other researchers to replicate our results. This way we can meet requirements related to reproducible research. Since user consent is not required, there is no consent bias and the findings are “generalizable” in that sense that the full population (i.e. all politicians that use Twitter) is included in our sample. Lastly, since the ethical implications of our Twitter studies are limited, there is no communication required with the data subjects or other instances. Our research is communicated in the scientific community. The ethical considerations are summarized in Figure 2.1.




Figure 2.1: Ethical considerations of the Twitter data collection and research (framework based on Greene et al. (2019)).

### 2.1.2 Facebook data (Part II)

Researchers that wish to access a user’s personal Facebook data need to obtain permission from Facebook as well as from the user himself. First, a thorough review process (called App Review)<sup>7</sup> needs to be passed, so that Facebook can make sure that private data of their customers are not misused. The submission should contain a detailed description of which data will be accessed through your application and how it will be used, including screen recordings of how the application will work. Facebook reviewers will review this submission and test your application to verify that it’s in compliance with their usage policies. Second, after being granted access to the Facebook API, researchers can ask users for permission to access items of their Facebook data via Facebook Login. Facebook Login<sup>8</sup> is a developers’ tool that allows users to authenticate using their Facebook credentials and log into an application. After authenticating, the user can provide permission to access certain items of their Facebook data (see a mockup example in Figure 2.2). Unfortunately,

<sup>7</sup> <https://developers.facebook.com/docs/facebook-login/review>

<sup>8</sup> <https://developers.facebook.com/docs/facebook-login/>

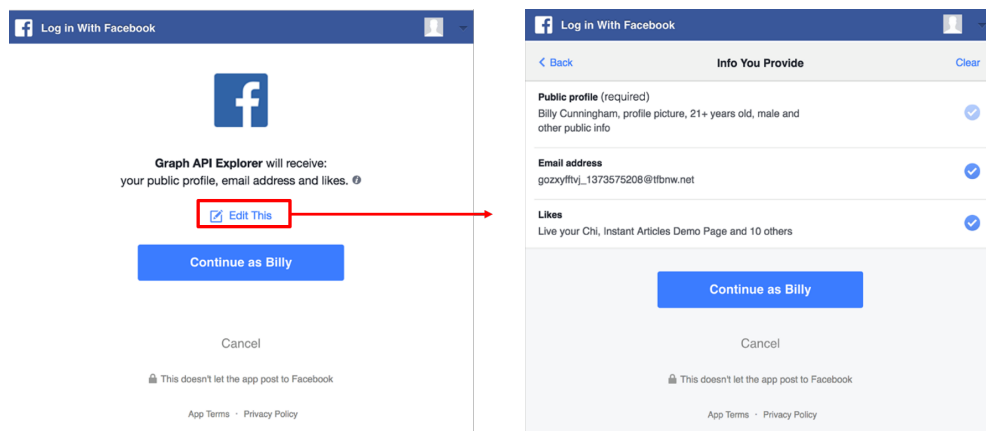


Figure 2.2: Mockup example of the Facebook Login button and the edit permissions window.

the data collection process described here is no longer available today. After the Cambridge Analytica scandal, the Facebook Platform has further restricted data access in August 2018, so that it is no longer allowed to access Facebook profile data for academic research –unless in collaboration with Facebook, e.g. through the Social Science One project.<sup>9</sup>

For the Belgian data, the data collection took place in two waves in March and June 2018. In the first wave (March 2018), a detailed survey with questions on socio-demographics, media consumption, political preference and attitudes was sent to a representative panel of around 4,500 respondents.<sup>10</sup> From these respondents, around 12% agreed to give us access to their Public Page Likes,<sup>11</sup> via Facebook Login. Based on the data of this small set of respondents, models were built to predict gender, age, political leaning and party preference based on Facebook Page Likes. These initial prediction models were used to develop a tool that shows participants which characteristics can be inferred about them based on their Facebook Likes. The goal of this tool was two-fold: First, to convince people to participate in our study and second, to create awareness about the personal information that you might disclose about yourself on Facebook. In the second wave (May–June 2018), this tool was disseminated through the online webpages of popular Flemish newspapers to reach a broad audience. Via an online webpage, users could give consent to collect their Facebook Likes and they were also asked to complete 12 survey questions about their media consumption and political preference. In return for completing the survey, participants could see a prediction of their gender, age, ideology and party

<sup>9</sup> <https://socialscience.one/>.

<sup>10</sup> More details on the survey questions will be provided in Chapter 5

<sup>11</sup> *Public Page Likes* are the public Facebook pages that a user has liked and that show up as being liked in the About section of that person's profile (see [https://www.facebook.com/help/171378103323792?helpref=uf\\_permalink](https://www.facebook.com/help/171378103323792?helpref=uf_permalink)). We do not collect likes or emotional reactions (Love, Haha, Wow, Sad, and Angry) to Facebook posts.



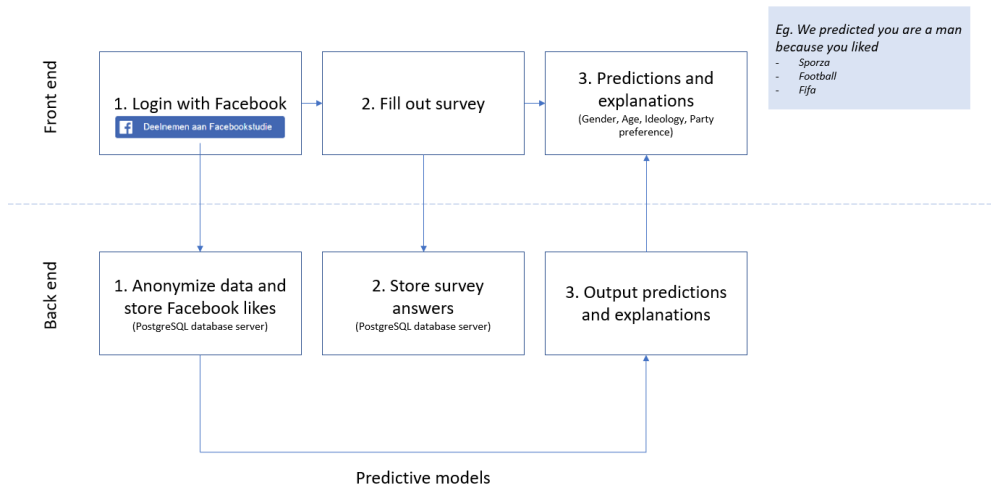


Figure 2.3: Overview of the front- and back end of the webpage that was created to collect Facebook like data for our study.

preference based on the Facebook Likes they provided. All data were anonymously and securely stored on a PostgreSQL database server. An overview of the front- and back end of our webpage can be found in Figure 2.3. After the data collection period, the data were downloaded from the server, stored in an encrypted file on a local computer, and the initial database was deleted.

Since the survey answers and Facebook Likes are private and highly sensitive data, the ethical implications reach further than with the Twitter data and deserve the necessary attention (see Figure 2.4). All data were collected with informed consent. We clearly stated what data would be gathered and for what purpose. Respondents were required to agree with our privacy policy to participate in our study. The statement includes information on which data will be collected, for what purpose and for how long it will be stored. Participants also had the right to stop their collaboration at any time and ask for their data to be removed. Furthermore, no other data than the data needed for our study were collected. The data are anonymized immediately, meaning that all personal identifiers (i.e. name and Facebook ID) are permanently removed. Yet, because of the risk of reidentification (Narayanan and Shmatikov, 2008), the –anonymized– data will under no circumstances be made public or shared with other instances. The data are stored locally and encrypted for optimal security. The data are used for scientific purposes only and will be removed permanently after ten years. Results are shown on an aggregate level only and will never contain individual examples, nor will we explicitly mention pages that were liked by less than 30 participants. Albeit absolutely necessary to protect privacy of our respondents, our privacy policy demonstrates two important implications for scientific generalizability. First, informed consent automatically

induces self-selection or consent bias, and generalizing our results to the total population is difficult. Second, since we can only share aggregated results, this complicates the reproducibility of our study. However, the complete computer code and aggregated data can be made publicly available to assess scientific quality. Finally, a last important aspect of ethical research conduct is clear communication with data subjects (participants) before and after the study. Before the start of our data collection we compiled a clear privacy statement, in collaboration with the university's data protection officer and with approval by the Ethics Committee for the Social Sciences and Humanities (EA SHW).<sup>12</sup> The full privacy statement was communicated with participants and can be found on the project website,<sup>13</sup> where we have also communicated a concise summary of our results.

Finally, the American Facebook data were collected in a very similar set-up by the Social Media and Political Participation (SMaPP) Lab at New York University. A panel survey on (social) media use was conducted during the 2016 U.S. presidential election ( $N = 3,500$ ) over three waves. Respondents completed a survey and were asked if they would be willing to supply information about their own Facebook activity. The data that was requested includes public profile information, Timeline posts (including text and links if available), Public Page Likes, and what Facebook saves as religious and political views. Via Facebook Login respondents could approve sharing all of the given types of information, selectively approve only some of these types of information, or approve none of them. All respondents who agreed to share information consented to a privacy policy that informed them that they could deactivate the application at any time, that no personally identifying information would be shared, and that this application will not access the profile information of any friends, groups, or other information associated with their profile page. The data are also anonymous and stored securely. The research design has been approved by the New York University Institutional Review Board (IRB-12-9058, IRB-FY2017-150).




Figure 2.4: Ethical considerations of the Facebook data collection and research (framework based on Greene et al. (2019)).

<sup>12</sup> <https://www.uantwerpen.be/en/research/management/quality-assurance/ethics-screening/eashw/>

<sup>13</sup> [https://www.uantwerpen.be/nl/projecten/nws-data/privacybeleid/\(Dutch only\)](https://www.uantwerpen.be/nl/projecten/nws-data/privacybeleid/(Dutch%20only))

## 2.2 Data mining concepts

This section provides a brief introduction to some essential data mining concepts, including classification techniques, dimensionality reduction, and performance metrics. These methods will be applied in the following studies.

### 2.2.1 Classification techniques

The goal of a classification model is to identify to which class an observation belongs. Based on a training set  $X$  with  $n$  instances or observations and target variable or class  $y$  for each instance, a classification model learns a function  $f(x) = \hat{y}$ , to predict the target variable for unseen instances. For example, based on the Facebook likes of  $n$  Facebook users and their ideological leaning  $y_i \in \{left, center, right\}$ , a classification model learns to predict ideological leaning for unseen users, based on their Facebook likes (see Figure 2.5). If the target variable consists of multiple classes (e.g. left, center, and right), this is called multi-class classification. In case of a binary target variable (e.g. left or not), this is binary classification.

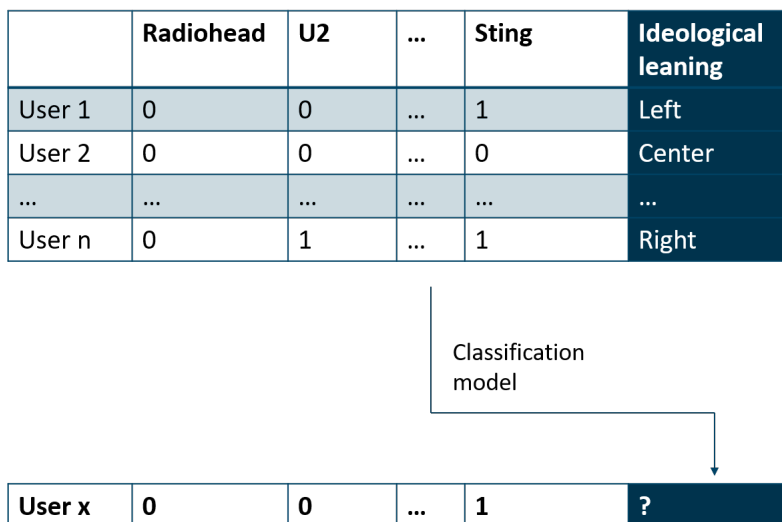


Figure 2.5: Fictitious example of a classification model that predicts ideological leaning based on Facebook Likes. The features consist of Facebook pages, and the value will be 1 if the user liked the page and 0 if not. The target value is ideological leaning and consists of three classes: left, center, and right.

Some commonly-used classification techniques include Logistic Regression (LR), Decision Trees, Support Vector Machines (SVM), Naive Bayes, and Neural Network

models. In this thesis, mainly Logistic Regression is used for two reasons: first, according to a large benchmark study (De Cnudde et al., 2017), it is the best performing technique in terms of predictive performance for behavioral and textual data (which was confirmed by our own experiments on the data as well) and second, the coefficients of an LR model are very intuitive to interpret. Interpretability of the model is a requirement to gain relevant insights from the data.

**LOGISTIC REGRESSION** With a binary response coded in the form  $Y \in \{0,1\}$ , logistic regression models the ratio of chances  $P(Y = 1|X = x)$  and  $P(Y = 0|X = x)$  (Tibshirani et al., 2015):

$$\log\left(\frac{P(Y = 1|X = \mathbf{x})}{P(Y = 0|X = \mathbf{x})}\right) = \mathbf{w}^T \mathbf{x} + b \quad (2.1)$$

where  $\mathbf{x}$  represents a  $d$ -dimensional input vector,  $b$  is an intercept term, and  $\mathbf{w}$  a vector of regression coefficients. The higher the coefficient of the variable in the model, the higher the positive association with the target. Alternatively, this can be written as the logistic (sigmoid) function (see Figure 2.6):

$$P(Y = 1|X = x) = \frac{1}{1 + e^{-\mathbf{w}^T \mathbf{x} + b}} \quad (2.2)$$




Figure 2.6: The sigmoid function maps the outcome of a linear function  $z = \mathbf{w}^T \mathbf{x} + b$  to the range  $[0, 1]$ . Source: Jurafsky and Martin (2018).

The unknown parameters  $\mathbf{w}$  and  $b$  of the linear model are estimated by maximizing the conditional likelihood that the predicted label for each training observation is equal to the true label. We do this by minimizing the distance between the predicted and true labels, and we call this distance the loss function or the cost function. If we use notation  $y_i \in \{-1, 1\}$ , we can define the loss function as:

$$\sum_{i=1}^N \log(1 + e^{-y_i(\mathbf{w}^T x_i + b)}) \quad (2.3)$$

In the case of logistic regression on sparse data, simply minimizing the standard loss function will lead to estimates that are not unique and can overfit the data. To avoid this, an extra constraint is introduced to the function that penalizes the weights of coefficients: regularization (Tibshirani et al., 2015). The most frequently-used regularization techniques are L1 regularization (Lasso regression) and L2 regularization (Ridge regression).

**L1 regularization** adds the sum of the absolute values of the coefficients as a penalty term to the loss function.

$$\min_{\mathbf{w}} \sum_{i=1}^N \log(1 + e^{-y_i(\mathbf{w}^T x_i + b)}) + \lambda |\mathbf{w}| \quad (2.4)$$

The regularization parameter  $\lambda$  models a trade-off between the complexity of the model and minimization of the prediction error. Higher values of  $\lambda$  will result in more regularization and thus a penalty for large feature weights or high complexity. The value of  $\lambda$  is optimized on a separate part of the training data, the validation set. L1 regularization zeroes out small coefficients, which results in natural feature selection.

**L2 regularization** adds the sum of the squared of the coefficients, and leads to small but non-zero weights:

$$\min_{\mathbf{w}} \sum_{i=1}^N \log(1 + e^{-y_i(\mathbf{w}^T x_i + b)}) + \lambda |\mathbf{w}|^2 \quad (2.5)$$

To find the optimal weights, the loss function can be minimized with stochastic gradient descent or with batch gradient descent (Jurafsky and Martin, 2018). We apply the *Scikit-learn* implementation of logistic regression with the default ‘lbfgs’ solver.<sup>14</sup>

<sup>14</sup> [https://scikit-learn.org/stable/modules/generated/sklearn.linear\\_model.LogisticRegression.html](https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LogisticRegression.html)




Figure 2.7: Illustration of matrix factorization of a matrix  $X$  consisting of  $n$  unique instances and  $m$  unique features into two non-negative matrices  $W$  and  $H$  of the original  $n$  instances by  $k$  dimensions and those same  $k$  dimensions by the  $m$  original features.

### 2.2.2 Dimensionality reduction

A useful method when working with high-dimensional data is dimensionality reduction. The idea is to represent the data in a lower dimensional space without too much loss of information. The original data matrix  $X_{n \times m}$  with  $n$  unique instances and  $m$  unique features is split into two matrices  $W_{n \times k}$  and  $H_{k \times m}$  such that:  $V \approx WH$  (see visual representation in Figure 2.7). The  $k$  columns of  $W$  are the dimensions, and each instance will have a representation in the new  $k$ -dimensional space. The matrix  $H$ , represents the relationship between the new dimensions and the original features (Clark and Provost, 2015).

This does not only speed up computations but might also improve the interpretability of the data. It is a popular technique in information retrieval because it groups related features together. When used in the context of text mining, we refer to this techniques as topic modeling. In that case, the term-document matrix is represented by two matrices, one containing the words per topic and one containing the topics per document (O’callaghan et al., 2015). The quality of the components depends on the choice of  $k$ , the number of dimensions. When  $k$  is too low, the components will be overly broad, while setting  $k$  too high can lead to many highly-similar components. Depending on what the goal of the dimensionality reduction is, different solutions are proposed to optimize the number of  $k$ . For insightful results, several values for  $k$  can be tried and the produced dimensions can be manually inspected for their coherence or scored by coherence metrics (Stevens et al., 2012). Albeit useful to discover hidden structures in the data, dimensionality reduction does not always improve final classification performance (Conover et al., 2011).

Several dimensionality reduction techniques exists, such as Principal Component Analysis (PCA), Singular Value Decomposition (SVD), Latent Dirichlet Allocation (LDA), and Non-negative Matrix Factorization (NMF). The latter has the nice prop-

erty that it produces feature vectors with only non-negative values, which facilitates the interpretation of the original data in terms of the latent factors (Contreras-Piña and Ríos, 2016; Lee and Seung, 1999). Therefore, NMF will be used repeatedly throughout this thesis to gain additional insights from the data.

**NMF** NMF is applied in multiple domains to decompose a non-negative matrix into two non-negative matrices (Sorzano et al., 2014):

$$X \approx WH \quad (2.6)$$

while minimizing the objective function:

$$J = \frac{1}{2} ||X - WH|| \quad (2.7)$$

Several algorithms exist for solving this optimization problem (Wang and Blei, 2011), we will make use of the *Scikit-learn* implementation based on coordinate descent.<sup>15</sup>

### 2.2.3 Performance metrics

Finally, we need to evaluate the performance of a classifier, or how well the classifier manages to assign the correct class to instances. As mentioned in Chapter 1, prediction performance is evaluated on a part of the data that was not used for training the model: the test data (also known as hold-out data). The test data contains known target labels, but is not used for training the model, which allows to assess the model's generalization performance on unseen data.

We'll consider binary classification where the target class can be positive or negative (e.g. left or not left). Typically, a classification method will assign an output *score* for a particular instance. This score reflects the probability of the instance belonging to the positive class. Subsequently, a threshold is applied to transform the score into a prediction regarding the label (positive or negative) of the instance.

**ACCURACY** Accuracy is a popular metric because it is very easy to measure and intuitive to understand. It is the number of correct classifications divided by the total number of classifications made (Provost and Fawcett, 2013).

$$Accuracy = \frac{\text{Number of correct classifications}}{\text{Total number of classifications}} \quad (2.8)$$

Unfortunately, this number is a bit too simplistic. If we have 90% negative instances in the data, a classifier that always predicts the negative class will achieve 90%

<sup>15</sup> <https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.NMF.html>

accuracy. This sounds very good at first sight, but of course the classifier does not perform well in distinguishing between positive and negative cases.

**CONFUSION MATRIX** To get a better insight in the missclassifications of the model, the actual and predicted labels can be summarized in a confusion matrix. For a problem involving  $n$  classes this is an  $n \times n$  matrix with the actual classes (columns) and the predicted classes (rows). For a binary classification the confusion matrix is a  $2 \times 2$  matrix, as shown in Table 2.1.

Table 2.1: Confusion matrix

	Actual positive	Actual negative
Predicted positive	True Positives (TP)	False Positive (FP)
Predicted negative	False Negative (FN)	True Negatives (TN)

A positive instance that is classified as a positive is a true positive, otherwise it is referred to as a false negative. A negative instance that is correctly classified as such is a true negative, otherwise it is a false positive. Based on the confusion matrix, a number of traditional evaluation metrics can be defined:

- Recall, or accuracy on the positive class (also called sensitivity or true positive rate TPR):

$$Recall = TP / (TP + FN)$$

- Precision, or the accuracy on the positive predictions:

$$Precision = TP / (TP + FP)$$

- Specificity, or the accuracy on the negative class (also called true negative rate TNR):

$$Specificity = TN / (FP + TN)$$

- F1-score, the harmonic mean between precision and recall:

$$F1 = 2 \times \frac{precision \times recall}{precision + recall}$$

**AUC** The evaluation metric we will use frequently throughout this work is Area Under the ROC Curve (AUC). The Receiver Operating Characteristics (ROC) curve



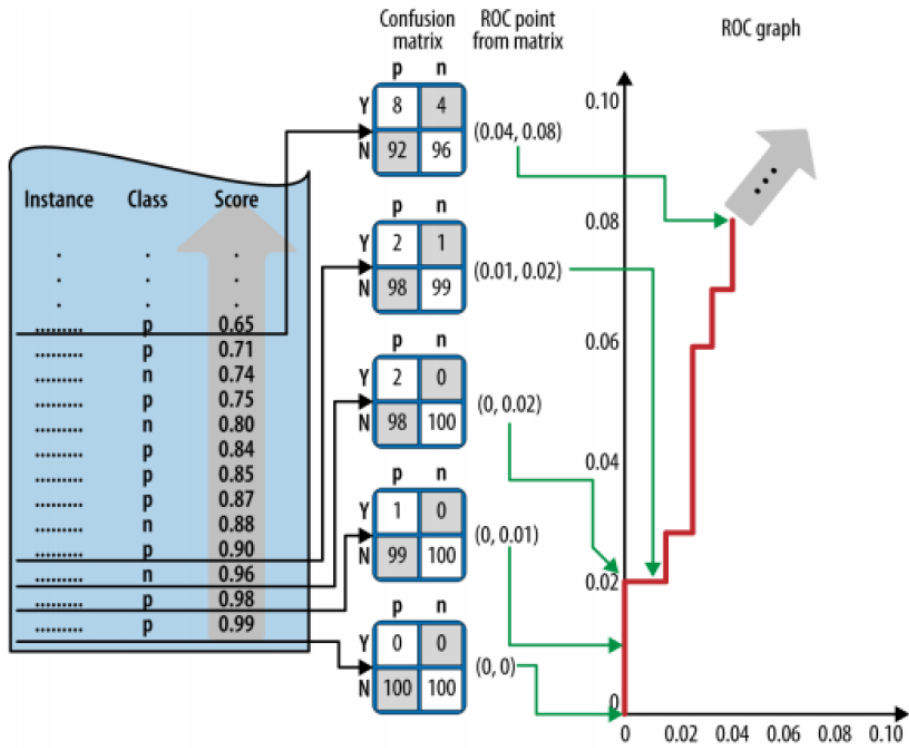


Figure 2.8: An illustration of how a ROC curve is constructed from a test set. The model assigns a score to each instance and the instances are ordered decreasing from bottom to top. Plot each FPR (x-axis) and TPR (y-axis) pair for descending thresholds. Source: Provost and Fawcett (2013)

plots the the true positive rate (TPR) of a classifier on the y-axis and their false positive rate (FPR)<sup>16</sup> on the x-axis (Provost and Fawcett, 2013). The plot is built by ranking the classifier's prediction scores in descending order and by ranging over all possible thresholds from high to low. The ROC-curve starts in the origin (0,0), where the threshold is at its highest so that all points are predicted negative, hence  $TPR = 0$  and  $FPR = 0$ . Subsequently, the threshold is lowered and the TPR and FPR for each value of the threshold are plotted. The ROC-curve ends in the point (1,1), where all instances are predicted positive. Ideally, the ROC-curve crosses the point (0,1) which means all instances are correctly predicted. Figure 2.8 illustrates the construction of the ROC-curve.

The Area Under the ROC Curve (AUC) summarizes the ROC curve in one scalar between 0 and 1. The more the ROC curve is located towards the upper left corner, the higher the AUC value will be. AUC can be interpreted as the probability

<sup>16</sup>  $FPR = 1 - TNR$

that the model ranks a random positive example higher than a random negative example (Flach et al., 2011). A perfect model would achieve an AUC of 100%, while an AUC of 50% is equivalent to the performance of a random model. AUC is a commonly-used metric in data science because it has a number of merits. It is a single number, allowing for straightforward comparison across several classifiers. It is an objective measure of the performance of a classification model, independent of the parameter choices such as the threshold, and it is unaffected by the frequency of the classes (Hand, 2009). However, since AUC summarizes the performance over the entire range of values of the classification threshold, it is unfit to inspect performance for a specific (set of) threshold(s) (Hand, 2009). Furthermore, AUC assumes misclassification costs are equal for both classes (misclassifying a positive as a negative is equally bad as misclassifying a negative as a positive). For many applications this is not the case. For example, a false positive cancer screening test might be considered less severe than a false negative test. In the former case, the patient might undergo some extra tests, while in the latter case the cancer formation goes unnoticed and untreated. When misclassification costs are important, AUC is not the most suitable metric to measure and compare classifier performance. For the applications in this thesis however, the goal is to measure how well a classifier discriminates between two classes independent of the threshold settings, and AUC is thus a perfectly fit measure.

## Part I

# TWITTER TO STUDY COMMUNICATION AND INTERACTIONS



## Issue communication on Twitter

---

Party competition in Western Europe is increasingly focused on “issue competition”, which is the selective emphasis on issues by parties. The aim of this study is to contribute methodologically to the increasing number of studies that deal with different aspects of parties’ issue competition and communication. We systematically compare the value and shortcomings of three exploratory text representation approaches to study the issue communication of parties on Twitter. More specifically, we analyze which issues separate the online communication of one party from that of the other parties and how consistent party communication is. Our analysis was performed on two years of Twitter data from six Belgian political parties, comprising of over 56,000 political tweets. The results indicate that our exploratory approach is useful to study how political parties profile themselves on Twitter and which strategies are at play. Second, our method allows to analyze communication of individual politicians, which contributes to classical literature on party unity and party discipline. A comparison of our three methods shows a clear trade-off between interpretability and discriminative power, where a combination of all three simultaneously provides the best insights.

### 3.1 Introduction

Issues and issue preferences form the raw matter of politics. The classic theory of democratic representation states that voters are expected to vote for parties that best represent the issues they deem important and that best represent their positional policy preferences on those issues (Thomassen and Schmitt, 1997). Therefore, parties try to steer the debate in the direction of the issue they have a strong profile or reputation on; this yields them an electoral advantage. Furthermore, the fragmentation of party landscapes across Europe in recent decades has increased the number of issues parties put forward. This explains why party competition in Western Europe has increasingly focused on the battle about which issues should dominate the party political agenda, i.e. “issue competition” (Green-Pedersen, 2007). The growing importance of issues in party politics, is also reflected by the rising attention for and proliferation of theories dealing with issue competition and communication (e.g. De Sio and Lachat, 2020).

Traditionally, research would examine party manifestos, campaign ads or press releases to study strategic issue communication choices (Tresch et al., 2017). However, nowadays social media represents an interesting alternative, as it is perhaps the most widely accessible form of party communication, with higher temporal adaptability and interaction potential (De Sio and Lachat, 2020). There is growing scholarly interest in parties’ issue communication and strategies on social media (Vargo et al., 2014; Van Dalen et al., 2015; Van Ditmars et al., 2020). However, the high volatility of social media communication in combinations with relatively short and less formal text complicates automatic coding methods and party-level analysis. Therefore, the main aim of this study is to contribute to the rapid increase of studies that deal with different aspects of parties’ issue communication on social media.

Especially Twitter is increasingly used by political parties and politicians to communicate with citizens, but even more so with opinion leaders and journalists (Jungherr, 2016; Vargo et al., 2014). We accept the press-release assumption of political parties on Twitter as suggested by De Sio and Lachat (2020) and extend this to individual politicians of a party. It states that, irrespective of the amount and type of followers a party’s Twitter account might have, parties use Twitter as a way to communicate messages to the media and the public, like a press release, even in countries with low or elite-only Twitter penetration (Kreiss, 2016; Parmelee and Bichard, 2011).

In this study, we contribute to the issue competition literature by analyzing the issue competition of Flemish political parties on Twitter. More specifically, we are interested in how political parties differentiate themselves issue-wise from other parties in a multi-party system. We specifically focus on the emphasis they put on issues and not on their position towards issues. For instance, the theory of issue ownership states that parties can “own” issues if they are considered by the voters at large as the “best” party to deal with the issue (Petrocik, 1996; Walgrave et al., 2015). Hence, it is in a party’s interest to make sure that the issues it owns are high on the

priority list of voters. That is why parties tend to focus on their owned issues in their communication. Although several studies confirm that parties indeed focus on their issues, others show that parties “trespass” frequently and also address issues owned by their competitors (Damore, 2005). According to the recently developed issue yield theory, parties are more flexible and (ideologically) free to address issues that are not associated with the party as long as the party has a policy position on the issue that matches the party and if that position is also widely shared in the general electorate (De Sio and Lachat, 2020). While issue ownership and issue yield theory expect differences in the issue communication of parties, issue salience theories stress that parties and politicians address the issues that are high on the public and/or media agenda. By surfing the waves of issues that dominate the news, politicians can attract media attention for their political work (Van Santen et al., 2015; Wagner and Meyer, 2014). If issue salience would dominate politicians’ communication—they all follow the salient issues—we would expect to find few differences in parties’ issue emphasis. Our method is not designed to capture these salient issues in party communication, but rather focuses on the differences in issue communication strategy. In other words, we focus on the distinctive part of issue communication rather than overall issue communication strategies.

Second, as we do not study the party as a single, united actor but rather study individual parliamentarians; we examine how consistent and coherent parties communicate about issues. Or, in other words, do politicians of the same party communicate about the same issues? Especially in election times a consistent issue strategy and clear, recognizable communication are valuable assets for persuading and retaining voters. Aligning online communication of all party representatives might be a beneficial strategy (Van Dalen et al., 2015). There also are reasons for politicians of the same party to address different issues. For instance, individual politicians may try to emphasize the issues they are specialized in to signal their expertise, and compete with politicians inside and outside their own party by emphasizing distinct issues (Peeters et al., 2019).

We propose an exploratory approach based on predictive modeling to find the most discriminative issues per party. The advantages of this exploratory approach are threefold. First, it allows researchers to move beyond an exclusive focus on frequency when analyzing issue communication. Rather than focusing on the most frequent issues per party (which could be similar for all parties), we argue it is more interesting to focus on the issues that differentiate one party from the others. Second, it does not require manual issue-coding of (a part of) the tweets, which is often labor-intensive and time-consuming. Third, an exploratory approach can contribute to existing theory by increasing our understanding of how parties try to profile themselves and which mechanisms and strategic choices drive issue communication. More specifically, per political party and based on the content of the tweet, a classification model is built to predict whether the author of the tweet belongs to the political party. We systematically compare three ways to represent the content of a tweet: (1) an expert-driven approach based on dictionaries, (2) a data-driven approach based

on a bag of words method, and (3) another data-driven approach based on topic modeling. Before we turn to explain our data collection and discuss our results, we summarize established text classification methods in the field of politics and motivate our alternative approach.

### 3.2 Automated content analysis

Grimmer and Stewart (2013) argue that the understanding of language to know what political actors are saying and writing is central to the study of politics. Yet, the sheer volume of existing political texts does not allow for the manual reading and interpretation of all these documents. Automated content methods, however, can make the systematic analysis of large-scale text collections possible. For content analysis of political texts, typically two methods are considered: dictionary methods, based on the relative frequency of predefined key words in a document and supervised learning methods where the algorithm learns to classify documents into categories using a labeled training set (Grimmer and Stewart, 2013). Typically, when one is interested in party-level issue communication, one would classify texts into policy issues using one of both approaches, and aggregate results to learn the frequency of communication per issue at the party level. Next, we discuss how both methods can be used for the automated classification of policy issues in texts; after which we will explain why focusing on issue frequency might not be optimal to study issue communication by political parties.

To define issues, political scientists around the world often refer to the Comparative Agendas Project (CAP) codebook, consisting of 21 major issues (e.g. Environment, Macroeconomics), and more than 200 sub-issues.<sup>1</sup> Sevenans et al. (2014) manually compiled a Dutch dictionary of indicator words for each of the 21 CAP issues and showed it performs relatively well for issue classification. An important limitation of dictionary methods is that they depend on the quality of the predefined keywords and that dictionaries are of limited length, meaning that dictionaries are unable to capture all possible words related to a certain issue. When working with short texts such as tweets, the probability for dictionary words to appear in such a short text is low (Zirn et al., 2016). Moreover, with new words or terms being generated, a dictionary—mostly designed for formal text—soon becomes outdated (Wu et al., 2018). At the same time, extending dictionaries to improve coverage might come at the expense of lower precision.

To overcome the drawbacks of dictionaries, supervised learning has become a popular alternative. With supervised learning, the relevant features of the text and their weights are automatically estimated from a labeled data set (Barberá et al., 2019). Often-used methods for text classification are Logistic Regression, Support Vector Machines and Naive Bayes (Paul et al., 2017). Also recently, different variations of neural networks have been proposed for text classification (Lai et al., 2015). A

<sup>1</sup> <http://www.comparativeagendas.net/pages/master-codebook>



notable challenge for the use of supervised learning, however, is that training a well-performing classifier requires a large training dataset coded by humans, where all policy issues of interest are well represented.

Annotating data is labor-intensive and several solutions have been proposed to reduce the coding work to a minimum; such as employing labeled data from a related task but different corpus, or using hashtags or well-defined keywords as annotations instead of human codings (Hasan et al., 2014; Gupta and Hewett, 2020). Next to that, semi-supervised learning (Van Engelen and Hoos, 2020) and transfer learning (Terechshenko et al., 2020) can be relevant to train a classifier when labeled data is scarce. The latter have been shown to outperform traditional classifiers with the same amount of (coded) training data (Terechshenko et al., 2020) but are increasingly complex and computationally demanding.

To sum up, achieving reliable document classification is hard, especially when considering a large number of classes. It requires compiling and/or updating dictionaries that are applicable to fast-evolving social media texts, or training a classifier on labeled data for which reaching sufficient accuracy is challenging to say the least. Moreover, remember from Chapter 1 that even with an accurate document classifier, the conclusions based on these results can be biased. The reason is that we try to optimize the classification of individual documents (tweets) in predefined categories (CAP issues), while the end goal is in fact to estimate the frequency or proportion of communication about a certain issue in a collection of documents (e.g. what percentage of tweets by NVA is about Macroeconomics). Unfortunately, even a well-performing individual classification model can be biased when the goal is to estimate category proportions. Suppose that all misclassifications happen in the same category, then the statistical bias in estimating the aggregate proportions could be very high (Hopkins and King, 2010). Methods exist to correct for this bias, or that give approximately unbiased estimates of category proportions directly, but they still require a sufficient set of labeled data (Hopkins and King, 2010).

Finally, we argue that frequency of communication about a certain issue is in most cases not the object of interest. If all parties talk a lot about a certain issue, it is not inherent to a particular party's communication strategy. Therefore, it is more insightful to learn which policy issues are specific to one party but not to the others. In other words, how political parties differentiate themselves issue-wise from other parties. To illustrate this, have a look at the results of a frequency-based dictionary approach in Table 3.1. For half of the parties (left and center) the most-frequently discussed issues are almost completely identical. With a focus on frequency of communication we cannot differentiate between the issue strategies of these parties, as they seem similar at first sight.

Therefore, we propose to focus on discriminative issues (issues which distinguish one party from the others). We classify individual tweets according to the 21 CAP topics, using a dictionary. Subsequently, we apply supervised learning to automatically

label the political party that authored the tweet. When learning this task, the machine will learn which features (policy issues) are relevant to a specific actor's (party) communication (Gentzkow et al., 2016). As discussed, this approach has the downside that results will be biased by the performance of the dictionary. Hence, we propose a data-driven approach, that eliminates the need to classify individual tweets according to the 21 CAP issues upfront. Based on textual features, tweets are directly classified to the political parties, and the machine learns which textual features are relevant. Subsequently, human coders or domain experts can analyze the relevant features and label them with policy issues, which significantly reduces the amount of work compared to labeling the original texts. The disadvantage of this data-driven approach is that it will be harder to draw conclusions on issue competition, as also other aspects of communication are taken into account. On the other hand, the exploratory nature of this approach can also be an advantage, as it provides a more fine-grained look into party communication. Figure 3.1 provides a schematic representation of a frequency-based dictionary approach and the alternative methods we propose.

Table 3.1: Most frequent CAP issues for Flemish parties on Twitter when applying a traditional dictionary approach.

Party	CAP issues
Groen	1. Transportation 2. Environment 3. Macroeconomics
Sp.a	1. Environment 2. Macroeconomics 3. Transportation
CD&V	1. Education 2. Transportation 3. Macroeconomics
Open VLD	1. International affairs 2. Education 3. Transportation
NVA	1. Immigration 2. Macroeconomics 3. International affairs
Vlaams Belang	1. Immigration 2. Law and crime 3. Government operations

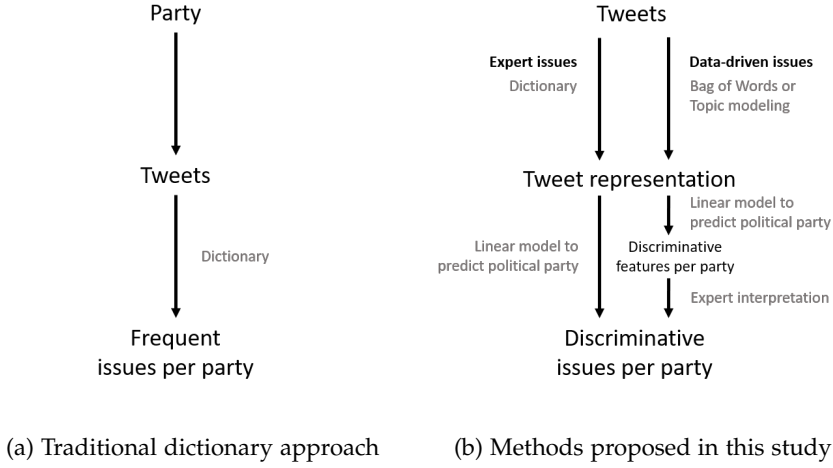


Figure 3.1: Schematic representation of a traditional dictionary approach (a) and the alternative methods we propose in this study (b).

### 3.3 Data and methods

In this study, we propose and validate the use of an exploratory approach to learn about issue communication and emphasis in Belgium (Flanders).<sup>2</sup> We have collected tweets from six Flemish political parties and their elected politicians. Per political party, we train a classification model that predicts whether the author of a tweet belongs to the political party or not, based on the representation—defined in three ways—of a tweet. The properties of the trained models are investigated to analyze issue communication per political party. First, the most discriminative features (with the highest coefficients in a linear model) show which issues distinguish parties' communication from one another (RQ1). In this study, we will focus on the top three most discriminative issues, but note that any other number can be chosen depending on the research desires. Second, the performance or discriminative power of the model per political party (measured by AUC, see Section 2.2.3) indicates how well the classification model can distinguish one party from the others. High discriminative power suggests that internal party communication is consistent and different from other parties (Gentzkow et al., 2016). Therefore, we will consider discriminative power per party as a proxy for internal consistency in party communication (RQ2). The research questions and method are summarized in Figure 3.2.

#### 3.3.1 Data collection

For a time period of two years between October 2017 and October 2019, we collected more than 256,000 tweets from the official Twitter accounts of the six political parties

<sup>2</sup> Replication code can be found on Github: [https://github.com/SPraet/issue\\_communication](https://github.com/SPraet/issue_communication)

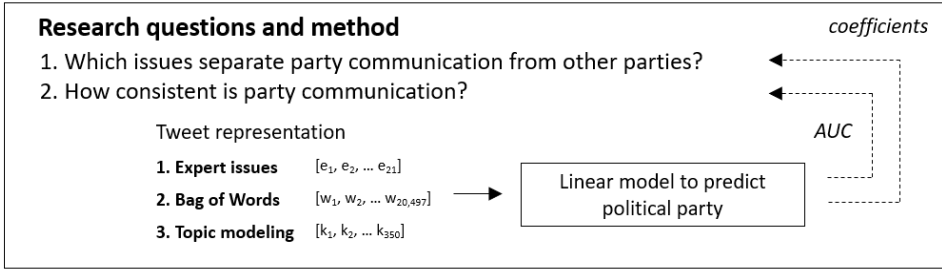


Figure 3.2: Overview of our exploratory approach to investigate issue communication by political parties on Twitter.

represented in the (Flemish and federal) parliament: the Greens (Groen), Social Democrats (Sp.a), Christian-Democrats (CD&V), Liberals (Open Vld), Flemish Nationalists (NVA) and the Radical Right (Vlaams Belang, VB) and all their elected party representatives in the national or regional parliament including cabinet ministers and party leaders. First, we only select original tweets from these accounts, i.e. we do not include replies or retweets. Next, we separate the issue tweets, namely tweets that deal with a policy issue, from the tweets that deal with private life or refer to non-issue related aspects of politics such as messages to announce a campaign rally. Removing private and non-issue related tweets results in a higher quality (less noise) dataset for our purpose. We use a trained classifier<sup>3</sup> to select the issue tweets, which results in a final dataset of around 56,000 tweets by 227 individual politicians and six political parties. The number of accounts and tweets per party can be found in Table 3.2.

Table 3.2: The number of accounts and tweets per party

Party	Number of accounts	Number of tweets
NVA (Flemish nationalists)	80	18,860
CD&V (Christian-democrats)	53	12,400
Open Vld (Liberals)	36	6,023
sp.a (Social-democrats)	31	6,545
Groen (Greens)	21	7,201
VB (Radical Right)	12	5,195

<sup>3</sup> An external classifier (<https://ccm.technology/>) was trained on more than 37,000 labeled Facebook posts of Flemish politicians, to distinguish between issue-related tweets, private tweets and non-issue related (campaign) tweets. To test the performance of this classifier on our dataset, a random subset of 500 tweets was selected and manually labeled. The accuracy of the classifier on this test set was 84% (11% false positives and 5% false negatives) and AUC was 92%. However, our approach is still applicable without this additional step and provides very similar results. The most predictive features per political party are largely the same and the relative predictive power is similar. Therefore, removing non-issue related tweets results in higher overall predictive power (because of less noise) but does not alter the conclusions

### 3.3.2 *Preprocessing of tweets*

Since the main interest of this research is to see how word usage in tweets might relate to political issues, we aim to reduce the event-specific information the tweets contain. Through intensive preprocessing we also want to reduce the noise that is common to social media texts (Han and Baldwin, 2011).

Tweets are first split into tokens and non-alphanumeric characters and stopwords<sup>4</sup> are removed. For Twitter specifically, this means that hashtags lose their ‘#-prefix and are handled as any other word. The use of user mentions, numbers and URLs in tweets is commonplace and might be informative for certain political issues; numbers playing an important role in financial news for example. However, we are not interested in the specific user, number or URL since it is unlikely that we can generalize from these. For that reason, these tokens are replaced with distinct placeholders.

Similarly, we argue that specific named entities (NE) in tweets are less informative to detect general policy issues. Using these words as features will cause our system to model specific events that occurred in the time-period of our data collection, rather than the more general policy issues that would be comparable to the expert dictionary. However, when it comes to named entities, the type of entity can still be informative for our purposes. Frequent mentioning of locations, for example, could be more indicative of issues like foreign affairs or defense, while frequent occurrence of organizations and products could relate to national economy. We use the Python library spaCy<sup>5</sup> for fine-grained tagging of named entities. We distinguish several types of named-entities such as locations, persons, organizations, products and events,<sup>6</sup> and replace them with their respective placeholders.<sup>7</sup> Lastly, we reduce word variation by lemmatizing the remaining tokens.<sup>8</sup> We are only interested in the lemma form of words because we aim to model their relatedness to political issues, regardless of their inflectional form.

### 3.3.3 *Tweet representation*

Before the actual modeling can start, the preprocessed tweets are transformed to a numerical representation. This will be done in three different ways, ranging from expert-driven to data-driven.

<sup>4</sup> We use the Dutch stopwords corpus from NLTK (<https://www.nltk.org/>).

<sup>5</sup> <https://spacy.io/>

<sup>6</sup> For a complete list of entity types, see <https://spacy.io/api/annotation#named-entities>

<sup>7</sup> To assess how named entities influence our results, we have also repeated the same experiments (as will be explained in the following sections) for the data with named entities included. These results indicate that it is indeed the case that we model very specific short-term events as well as names of party representatives etc. Though the results are -as expected- better in terms of classification performance (AUC), they provide little insight in the general political issues of party communication

<sup>8</sup> We used the pattern.nl module developed by CLiPS: <https://github.com/clips/pattern>

Table 3.3: Overview of the 21 CAP issues (Sevenans et al., 2014).

Code	Issue
t100	Macroeconomics
t200	Human rights
t300	Health
t400	Agriculture
t500	Labor and employment
t600	Education
t700	Environment
t800	Energy
t900	Immigration
t1000	Transportation
t1200	Law and crime
t1300	Social welfare
t1400	Community development and housing
t1500	Banking, finance and domestic commerce
t1600	Defense
t1700	Space, science, technology and communications
t1800	Foreign trade
t1900	International affairs and foreign aid
t2000	Government operations
t2100	Public lands and water management
t2300	Culture and arts

**EXPERT ISSUES** In the first method, we will use the Dutch CAP dictionary compiled by Sevenans et al. (2014) to transform every tweet in our collection to 21 CAP issues. More specifically, every tweet is transformed to a binary vector of length 21, where each value represents the presence of a CAP issue in the tweet (1 if the issue is present in the tweet and 0 if not). Multiple issues can be present in one tweet. Consequently, predictive models are built on this representation to predict to which of the six parties the tweet belongs.

To evaluate the performance of the CAP dictionary, a random subset of 9,280 tweets was manually coded for the 21 CAP issues.<sup>9</sup> We first separate political tweets from non-political tweets<sup>10</sup> and then apply the CAP dictionary to code issues. We experimentally found that the CAP dictionary provides the best results when assigning an issue to a text as soon as one relevant dictionary word appears in

<sup>9</sup> The tweets were coded by two coders who agreed in 44% of the cases on all labels. A more detailed overview of intercoder reliability per issue can be found in Table 9.1.

<sup>10</sup> Again, we apply the external classifier described before. The number of political tweets is 4954, or 54% of the evaluation set.

the text, in which case the accuracy<sup>11</sup> of the CAP dictionary is 35%, recall is 20% and precision is 63%.<sup>12</sup> The low recall of the dictionary resulted in many zero input features (only 24% of the tweets could be assigned at least one issue, see Appendix 9.1.1). Since the performance of the CAP dictionary on our tweets is low<sup>13</sup>, we introduce two data-driven approaches below.

**BAG OF WORDS** A first data-driven representation is a basic Bag of Words (BoW) approach, where each unique word corresponds to an input feature for the classification model.<sup>14</sup> This is still among the most commonly utilized methods in text classification (Barberá et al., 2019; Dun et al., 2020). Words are transformed into a numerical matrix using term frequency-inverse document frequency (tf-idf). The tf-idf matrix is used as input to predict to which of the six parties the tweet belongs. Afterwards, the most discriminative words will be manually interpreted in terms of the 21 CAP issues (see Section 3.3.5).

**TOPIC MODELING** Alternatively, feature construction can be done using topic modeling techniques (see Section 2.2.2). The idea is to extract latent topics from the collection of tweets, where each topic is a multinomial distribution over words, and to represent each tweet as a mixture of these topics (Chang et al., 2009). We will apply Non-negative Matrix Factorization (NMF)<sup>15</sup> to automatically extract topics from the political tweets. The NMF topics are learned from the collection of political tweets,<sup>16</sup> and the original tweets are represented by  $k$  topics. Next, classification models are built on this representation. We optimize the number of topics ( $k$ ) based on the performance of the subsequent supervised task: classification to one of the six parties. This way, the number of topics is set to 350, which is considerably higher than the 21 expert issues. Our data-driven topics are thus much more specific than the expert issues. Again, these data-generated topics will be manually interpreted in terms of the 21 CAP issues (see Section 3.3.5).

<sup>11</sup> Since this is a multi-label problem, *accuracy* refers to the percentage of tweets for which all labels were classified correctly.

<sup>12</sup> A more detailed evaluation per issue can be found in Appendix 9.1.1

<sup>13</sup> We tried to improve the performance of the dictionary by extending it with word embeddings (see Appendix 9.1.2). Although this results in higher recall; precision and accuracy are much lower.

<sup>14</sup> Including n-grams did not improve performance of the models, nor interpretation of the results. In fact, n-grams hardly were included in the most predictive features, and when they did it was in combination with a named entity, e.g. "ORG URL" or "says MENTION".

<sup>15</sup> We have also experimented with another technique: Latent Dirichlet Allocation (LDA) (Blei and Lafferty, 2006). In our setting however, the predictive models based on the topics produced with NMF achieve higher discriminative power than with LDA, which is why we will report the results using NMF topics

<sup>16</sup> Additionally, we tried to build the NMF topics on a larger background collection, including tweets from all Flemish media channels and political journalists. It did not lead to more interpretable or more accurate results than topic detection on the political tweets only.

### 3.3.4 *Classification models*

Per political party, a classification model is built to predict whether the author of the tweet belongs to the political party or not, based on the representation of the tweet (see Figure 3.2). From these models, we want to analyze the most discriminative features for each of the six parties. For this reason we choose to work with Logistic Regression with L2 regularization<sup>17</sup>, since the coefficients of this model are straightforward to interpret. Moreover, the discriminative power of this model showed higher or similar to the other classifiers in our benchmark<sup>18</sup> for the three different tweet representations. The coefficients and discriminative power of the trained models are investigated to draw conclusions on issue communication per political party.

### 3.3.5 *Evaluation*

We will systematically compare the three tweet representations in function of two evaluation criteria: discriminative power, or the ability to discriminate between political parties, and interpretability. First, to report the *discriminative power* of each model the last 20% of the tweets in our dataset are used as a separate out-of-time holdout set. We use the Area Under the ROC Curve (AUC)<sup>19</sup> to measure how well the trained models can classify the political parties based on the tweet representations. We calculate the weighted average AUC for the six classification models (one for each political party) to evaluate the discriminative power of our three different methods.<sup>20</sup>

Second, we define *interpretability* as the extent to which the most discriminative features correspond with the 21 CAP issues. When using the expert issue representation, the three most discriminative features are CAP issues and therefore by definition 100% interpretable. For the BoW and topic modeling representations we ask two independent domain experts to manually label the most discriminative features of the classification models with CAP issues (see example in Appendix 9.1.3). Usually, topics extracted by a topic model are interpreted by humans by looking at the

<sup>17</sup> More specifically, we use the scikit-learn implementation for logistic regression (Pedregosa et al., 2011a). The model parameters are optimized (for AUC) using 5-fold out-of-time cross validation: the training data is split in 5 folds, where first the 5th fold is used as a validation set while the previous folds are used for training, then the 4th fold is used for validation and the previous folds for training, etc. The regularization parameter (C) is optimized in the interval [0.001, 0.01, 0.1, 1, 10]. For the topic modeling representation, we first optimize the number of topics k, which ranges from 0 to 400 with a stepsize of 50 and then we optimize the regularization parameter C for the optimal k.

<sup>18</sup> Other classifiers in our benchmark include (Multilayer) Perceptron, Lasso Regression, Linear Regression, Support Vector Machine, Naive Bayes, Decion Tree and Random Forrest

<sup>19</sup> See Section 2.2.3 for an explanation of this evaluation metric

<sup>20</sup> Note that the weighted average AUC is used to compare the discriminative power of our three methods, while the AUC per political party is used to investigate consistency of party communication (see Figure 3.2).



top-weighted words per topic (Chang et al., 2009). We will look at the top 15 words<sup>21</sup> to assign a CAP issue to an NMF topic. Similarly, for the BoW we will assume that 15 words represent one CAP issue. Since we want to report the three most discriminative issues (see Section 3.3), we will show 45 words. We repeat the same experiment with different domain experts and a different set of most discriminative features from a model trained on a random subsample of the data. The average percentage agreement of the two experts is used as a measure for interpretability (referred to as *INT*).

### 3.4 Results

In the following sections we provide our results regarding the two questions we introduced earlier: (1) which issues separate the communication of parties from each other and (2) how consistent is party communication? The first question is answered by looking at the top three most discriminative issues per party. Additionally, we explore to what extent this issue communication is in line with existing theory on issue competition. The discriminative power of the model per political party provides us with an answer to the second question. A high discriminative power indicates that communication is coherent and consistent across individual politicians of the same party, while being distinct from other parties. Before we answer these questions, we will start with an evaluation of our three tweet representations.

#### 3.4.1 Comparison of tweet representations

The classification models are built on tweet representations defined in three different ways: expert issues, BoW and topic modeling (NMF). When comparing these three approaches, a trade-off between classification performance of the classifiers and interpretability of the features becomes apparent. With the BoW representation the classification models are best able to distinguish between parties, while the expert issues offer the most direct interpretation of policy issues (Figure 3.3). The topic modeling representation seems to balance both criteria.

The models based on expert issues have an average AUC of 59% meaning they are only slightly better at discriminating between parties than random. One explanation is the limited performance of the CAP dictionary when converting tweets to the expert issues. Additionally, even with a perfectly accurate dictionary, valuable information (e.g. specific word usage) is lost when reducing the tweets to 21 issues, and we cannot discriminate between different sub-themes within the same issue. On the other hand, results are 100% interpretable as the issues are constructed top-down from the CAP dictionary itself.

<sup>21</sup> Usually between 6 to 30 words are considered, so other options are possible as well.




Figure 3.3: A comparison of our three methods on both evaluation criteria shows a clear trade-off between interpretability and discriminative power.

With an average weighted AUC of 79%, the models based on BoW perform best at distinguishing between parties. The 45 most discriminative words are matched to the three most corresponding CAP issues (See Appendix 9.1.3 or one example in Table 3.4). This task is hard for domain experts since the most discriminative words are not necessarily thematically related, and therefore the average weighted interpretability is only 48%.

Table 3.4: The most discriminative features for the extreme right party (Vlaams Belang) when using the BoW approach, and the three most related CAP issues.

Party	Most discriminative features	CAP issues
VB	immigration, tomvangriek, islamization, vlaparl, immigration pact, mass immigration, islam, alien, immigration stop, immigrant, mosque, cordon, mosque, community, population, illegal, immigration policy, asylum seeker, multicultural, border, flanders ours again, concerning, URL, real, scum, immigrant, cause, country, people, people, terrorist, stop immigration, liberty, independence, our people first, protect our people, muslim, headscarf, so-called, government, even, elite, pact, madness	<ol style="list-style-type: none"> <li>1. Immigration</li> <li>2. Government operations</li> <li>3. /</li> </ol>

The discriminative power of the models based on the topic modeling representation (AUC = 68%) is higher than with the expert issues but lower than BoW. Per party we look at the three most discriminative NMF topics (each represented by 15 words) and manually assign the most corresponding CAP issue (See Appendix 9.1.3 or one example in Table 3.5). The expert interpretability is 84%, which indicates that domain experts mostly agree on which CAP issue corresponds to the NMF topic. This approach seems to find the best balance between discriminative power and interpretability.

Table 3.5: The most discriminative features for the extreme right party (Vlaams Belang) when using the topic modeling representation, and their corresponding CAP issues.

Party	Most discriminative features	CAP issues
VB	1. URL, action, and, due, youngsters, again, worry, ready, drawing, petition, life, share, right, thanks to, helping 2. country, border, safe, criminal, population, origin, illegal, deportation, alien, greatest, when, migrant, deport, hard, nationality 3. our, community, protect, security, proposals, economy, society, values, welfare, and, earn, pride, norm, farmer, resolution	1. Human rights 2. Immigration 3. Social welfare

#### 3.4.2 Which issues separate party communication from other parties?

For every party, the most discriminative issues are shown in Table 3.6. For the more extreme parties on both sides of the political spectrum, the three methods give consistent results. For the greens (Groen), that started as a one-issue party, the issue focus on the Environment is still irrefutable, while radical right politicians (Vlaams Belang) have a clear focus on Immigration. These results are in line with issue ownership theory,<sup>22</sup> stating that focusing on a few policy issues on which they have built a reputation is an effective strategy for parties to garner more votes. Another party that has a clear issue focus, at least partly in line with the issue ownership theory is, according to the different methods, is the NVA. Although the Flemish nationalists were traditionally not strongly focused on Immigration, in recent years they tried to “steal” the issue from the extreme-right party Vlaams Belang, which is also reflected in their communication on Twitter.

For the three traditional parties who are more situated in the center the issue focus is slightly more diffuse. The social-democrats of the Sp.a are linked to one of their core issues (Social welfare), but more often to an issue of a competitor (Environment, the core issue of the Green party). The Christen-democrats (CD&V) most often communicate on Education, an issue that is traditionally linked to the many catholic schools in the country and for which the cabinet minister is a leading figure of their party. The (economic) liberals (Open Vld) seem to communicate least consistent on the issues they own (Macroeconomics), although several issues have an economic dimension (e.g. foreign trade, banking).

In sum, many parties’ communication on Twitter is in line with the theory of issue ownership. For all parties, we find at least one issue that can be considered as an “owned” issue (see issues in bold in Table 3.6). However, most parties also seem to “trespass” their owned issues, in line with other issue competition theories. For example, the issue International Affairs is not owned by the liberal party Open Vld

<sup>22</sup> For issue ownership in Flanders, we rely on the study of Peeters et al. (2019) who asked Flemish respondents which party they instinctively thought about when hearing a certain issue. We consider an issue owned by the party if the percentage of respondents that linked a certain party with the issue is higher than 20%.

but they do have a minister for development cooperation in the federal government, which might be the reason for this specific issue focus. The reason opposition parties go beyond their owned issues is that they communicate about issues in reaction to what the government does. For example, the issue Defense is not owned by the socialist party Sp.a but in the period of data collection they heavily criticized the government decision to buy fighter planes. Finally, issue salience theory suggests that parties also respond to policy issues that are high on the public agenda (Van Santen et al., 2015; Wagner and Meyer, 2014). During the period of analysis these issues were Environment and Immigration. While concerns about the environment, and climate change in particular, were increasingly picked up by parties other than the Greens, the theme of immigration remained almost exclusively in the hands of the (radical) right. The data-driven methods allow to investigate sub-issues within issues, although this was not the focus of our study. For example, with respect to the salient issue of Environment, the Greens talk about a general climate policy, while the social-democrats and liberal party merely mention deposits on cans and small bottles, the Christen-democrats refer to their own important theme, namely quality of life, and finally, the Flemish nationalists discuss the efficiency of nuclear power plants driven by their approach of “eco-realism”.

Table 3.6: The CAP issues Flemish party representatives communicate about on Twitter.

Party	Expert issues	Bag of Words	Topic modeling
Groen	1. <b>Environment</b> 2. Transportation 3. Agriculture	1. <b>Environment</b> 2. / 3. /	1. <b>Environment</b> 2. / 3. /
Sp.a	1. Defense 2. Environment 3. Health	1. <b>Social welfare</b> 2. Environment 3. Macroeconomics	1. Environment 2. Government operations 3. <b>Social welfare</b>
CD&V	1. <b>Education</b> 2. Foreign trade 3. <b>Social welfare</b>	1. <b>Social welfare</b> 2. Transportation 3. <b>Education</b>	1. Environment 2. / 3. <b>Education</b>
Open VLD	1. Foreign trade 2. Banking and finance 3. Agriculture	1. International affairs 2. <b>Macroeconomics</b> 3. Banking and finance	1. International affairs 2. Environment 3. Immigration
NVA	1. Public lands and water 2. <b>Immigration</b> 3. Science and technology	1. <b>Immigration</b> 2. Government operations 3. <b>Law and crime</b>	1. <b>Immigration</b> 2. Energy 3. <b>Immigration</b>
Vlaams Belang	1. <b>Immigration</b> 2. Government operations 3. Human rights	1. <b>Immigration</b> 2. Government operations 3. /	1. Human rights 2. <b>Immigration</b> 3. Social welfare

Note: Issues printed in bold are owned by the party (Peeters et al., 2019). If none of the CAP issues matches with the set of words this is indicated with /.

### 3.4.3 *How consistent is party communication?*

To assess how consistent parties communicate we explore the discriminative power of the models per party (see Table 3.7). We assume that high AUC indicates consistent communication by the politicians of the considered party. For our three methods, the radical right party Vlaams Belang, is most consistent in their communication. This is partially due to the fact that this party pursues a clear positioning and association with one policy issue (Immigration). In addition, the lower number of party representatives is of course another explanation for more coherent communication. In that sense, it is remarkable that the N-VA, by far the biggest party with 80 representatives, scores not much lower in terms of consistency. This might be partly due to the high internal party discipline that characterizes Belgian parties (Depauw and Martin, 2009), and the N-VA in particular (Van Erkel et al., 2014). For all parties, AUC is higher for the data-driven methods than for the expert issues. This could indicate that party communication is more complex and not reducible to predefined issues. Indeed, with topic modeling we discover other characteristics of party communication rather than the policy issues they talk about. For example, one of the NMF topics for the liberal party (Open Vld) consists of English words (all other topics are in Dutch) and was apparently discriminative for Open Vld as it is the only party that occasionally tweets in English. Next to that, we often see party campaign slogans or hashtags among the most discriminative words, which can of course not be directly related to a policy issue.

Table 3.7: Classification performance and interpretability of the expert issues, Bag of Words and topic modeling representation.

	Expert issues		Bag of Words		Topic modeling	
	AUC	INT	AUC	INT	AUC	INT
Groen	60%	100%	82%	33%	71%	100%
sp.a	63%	100%	76%	50%	63%	67%
CD&V	57%	100%	81%	67%	70%	50%
Open Vld	61%	100%	79%	33%	71%	100%
NVA	56%	100%	76%	50%	66%	83%
VB	68%	100%	87%	33%	72%	67%
Weighted average	59%	100%	79%	48%	68%	76%

## 3.5 Conclusion and future research

Using three different tweet representations, we looked at which policy issues separate political parties on Twitter. Overall, our methods are remarkably good in distinguishing parties based on their (issue) communication. According to our results, especially the more extreme parties communicate clearly about the issues they “own”. This finding is in line with issue ownership theory which suggests

that political parties compete by raising attention for those policy issue that are positively associated with their party. On the other hand, several parties, mainly those in government, seem to trespass and also communicate about other issues, in line with other issue competition theories, such as issue salience or individual issue specialization and ministerial competences. The results indicate that our exploratory approach is useful to study how political parties distinguish themselves on Twitter and which strategies are at play. In addition, from the examination of the most discriminative words it becomes clear that a large part of communication on Twitter is event-driven, with parties talking about and reacting to current events that are limited in time. A more detailed temporal analysis could shed light on to what extend parties try and are successful to link these events to their owned issues.

By looking at the discriminative power of our models per political party we can draw conclusions about the consistency of communication by party representatives. This is highest for the more extreme (and also smaller) parties. Twitter is a much more personal communication channel than manifestos or press releases and individual politicians are free to tweet what they want (Peeters et al., 2019). Yet, for some political parties a classification model performs rather well in identifying their tweets based on the text only. As suggested by Gentzkow et al. (2016) the ease with which a machine learning model can infer a politician's party from their (written) language could be a measure for partisanship. A common language can be a key factor in creating group identity and party cohesion, but it can also increase inter-party hostility. An interesting direction for future research might be to look into how aligned all party representatives are in their communication, and to investigate communication strategy and its link to party composition (number, popularity, seniority, etc.) to explain the differences. This could be a useful contribution to the classic literature on party unity and party discipline that so far has not included the communication of individual politicians in their work (e.g. Depauw and Martin, 2009; Andeweg and Thomassen, 2011).

Lastly, with respect to our methodology, we think there's value in focusing on the distinctive character rather than just the frequency of communication. Classification models can distinguish one party from all other parties based on its communication, but they could also be applied to discriminate between two parties of interest (e.g. what is the difference in communication strategy of two nationalist parties NVA and Vlaams Belang). The expert- and data-driven approaches each have their advantages and disadvantages but by applying them simultaneously, different and complementary insights can be gained. The expert issues are insightful at the general issue level, but, next to being a result of low dictionary performance, the low AUC suggests that a lot of information is lost by trying to reduce political communication on Twitter to predefined issues. The low AUC could also suggests that political parties do not particularly differentiate themselves from their competitors in terms of issues but more in terms of specific content, as suggested by the higher AUC of the data-driven approaches. The data-driven approaches offer much more fine-grained insights at the event and even stylistic level of communication, at the expense of

interpretability at the issue level. Moreover, the data-driven approaches allow to analyze sub-themes within issues. Although this was not the main focus of our study, our methods could help to study issues at a more fine-grained level. Additionally, the results could even help to improve issue dictionaries by bringing forward synonyms or other related terms. For example, the herbicide “glyphosate” was topic for debate during the time period of analysis. The term is not included in the current CAP dictionary, but is clearly related to the issue “Environment”.

The methodology we propose is applicable to other (social media) text data and research questions as well. The expert-driven approach would benefit from improvements in document classification techniques. Recent advances in data-enhanced dictionaries, deep learning, transfer learning and semi-supervised learning offer exciting avenues for political text classification while at the same time introducing a lot of additional complexity and requiring ever more computing power. Adapting text classification to the volatility of social media remains a delicate exercise. Therefore, a promising method to study issue communication on social media is to start from a data-driven approach and use domain knowledge to interpret and understand the results.





## Parliamentary Twitter networks

---

Social media networks have revolutionized social science research. Yet, a lack of comparative empirical analysis of these networks leave social scientists with little knowledge on the role that contextual factors play in the formation of social relations. In this study we perform a large-scale comparison of parliamentary Twitter networks in 12 countries to improve our understanding of the influence of the country's democratic system on network behavior and elite polarization. One year of Twitter data was collected from all members of the parliament and government in these countries, which resulted in around two million tweets by almost 6,000 politicians. Social network analysis of the Twitter interactions indicates that consensual democracies are characterized by more dense parliamentary relations but also higher hierarchy and fragmentation compared to majoritarian systems. Secondly, parliaments with a high effective number of parties are more cooperative, which results in higher inter-party relations. Next to that, we show differences in the followers, mentions, and retweets networks that hold across all countries and political systems. Our empirical results correspond to established theoretical insights and highlight the relevance of institutional context as well as the platform characteristics when conducting social media research. With this research we demonstrate the importance and the opportunities of social network analysis for comparative research.

#### 4.1 Introduction

Social media has drastically changed the way people all over the world interact and communicate. Politicians are no exception. Today, social media is used as a new way to communicate and engage with voters, media and other politicians (Jungherr, 2016; Vargo et al., 2014). Especially Twitter is increasingly used by political parties and politicians to engage in political debate, publicly show support or disapproval, and communicate with other representatives (Teernstra et al., 2018). This new way of communication challenges some of the established theoretical insights in political science and introduces a number of technical obstacles. Simultaneously, it offers ample new opportunities to reassess how politicians interact with others. Network theory has been applied successfully to Twitter networks to offer insights in political polarization (Conover et al., 2011; Esteve Del Valle and Borge Bravo, 2018), opinion leadership (Borge Bravo and Esteve Del Valle, 2017), the underlying structure of political groups and countries (Cherepnalkoski and Mozetič, 2016), engagement with the public (Grant et al., 2010), etc. However, up to now, these studies have mostly focused on one country and do not allow for structured comparison across multiple countries to gain insights in contextual variables and country characteristics (Siegel, 2011).

Therefore, the aim of this study is to perform a large-scale comparison of Twitter networks to investigate the influence of institutional context on parliamentary relations. From September 2018 to September 2019, one year of Twitter data was collected from all members of the parliament and government in 12 countries with different political systems. This resulted in around two million tweets by almost 6,000 politicians.

With social network analysis and visualization we aim to explore three broad research subjects. First, we investigate whether the network properties of parliamentary Twitter networks are associated with the democratic system and functioning of the countries. We characterize the topology of the networks based on four widely-used network metrics: density, centralization, modularity, and the fraction of isolated users. We apply hierarchical clustering analysis to learn which countries are more similar based on their Twitter network properties, and link this to the electoral and party system of the countries. Secondly, we analyze inter-party communication as a measure of elite polarization along party lines. Next to linking this to the electoral and party system of the country, we also explore the correlation with ideological distance between parties. Lastly, we compare results across the followers, mentions and retweets network to learn how political interactions differ depending on the platform layer. To motivate why such comparative social network analysis can be valuable to improve our understanding of online social phenomena, we focus on the concept of elite polarization.

## 4.2 Parliamentary Twitter networks and elite polarization

A certain degree of political competition is necessary for a democratic system. Competing alternatives of public policy need to be presented to the public so that they can participate in the decision-making process. However, too much competition can lead to polarization which has detrimental effects on public decision making, as it stimulates partisan motivated reasoning, instead of decision making that relies on substantive arguments (Druckman et al., 2013). Therefore, political polarization, and the factors influencing it, have long been a central topic for political science.

The increasing popularity and use of social media have triggered debates about the effect of social platforms on polarization. Some claim that social media usage leads to increased polarization because individuals are more likely to engage with views similar to their own (Bimber and Davis, 2003). Bakshy et al. (2015) study 10 million Facebook users in the United States and observe that individuals are more likely to be exposed to information from like-minded individuals. Several studies suggest that political Twitter networks in the U.S. exhibit a highly segregated partisan structure (Conover et al., 2011, 2012; Barberá, 2015). Also in other countries strongly polarized structures have been observed on social platforms, including Switzerland (Garcia et al., 2015), Canada (Gruzd and Roy, 2014), and Italy (Quattrociocchi et al., 2016). In contrast, others argue that social media decreases polarization by exposing individuals to ideologically diverse information (Guess et al., 2018). Barberá et al. (2015) conclude that previous work may have overestimated the degree of ideological segregation in social-media usage in the United States. They find (especially liberal) individuals to engage in cross-ideological dissemination. Similarly, Boxell et al. (2017) demonstrate that greater internet use is not associated with faster growth in political polarization. In Europe, Moeller et al. (2018) do not find empirical evidence of increased polarization in the Netherlands, and Vaccari et al. (2016) indicate that cross-cutting interactions in Italy and Germany are less exceptional than expected.

In this research we focus on elite polarization from a network perspective, by analyzing the relational networks between parliamentarians on Twitter. Elite polarization can have important consequences for democracy. Polarization among the elite may influence mass polarization (Druckman et al., 2013), while elite bargaining and interaction are conducive to a stable democracy (O'donnell and Schmitter, 2013). Higley and Burton (1989) state that a liberal democracy is impossible without a "consensually united" national elite, which is characterized by dense and interlocked networks of communication and influence among the elite. Although consensual unity is a broader concept than polarization, a higher degree of polarization decreases the likelihood of consensual unity. In the first part of the analysis we will focus on the concept of consensual unity, which we will operationalize using network properties. In the second part we approach polarization more narrowly as the absence of interactions between opposing political groups (Conover et al., 2011).

The Twitter platform is well-suited to investigate interactions between parliamentarians. Twitter use among politicians is higher than among the general public, and also different: politicians mainly use it for political purposes, while citizens use it for political and non-political goals (Esteve Del Valle and Borge Bravo, 2018). Twitter's open character could foster more dialogue along ideological lines without party restrictions, or confine parliamentarians to partisan divisions similar to the offline world. For example, (Cook, 2016) find that the legislators' social connections on Twitter are less partisan than offline relations such as voting and co-sponsorship. On the other hand, Swiss politicians show a very strongly polarized structure in online support networks (Garcia et al., 2015). Similarly, communication flows of Catalan parliamentarians are found to be polarized along party and ideological lines (Esteve Del Valle and Borge Bravo, 2018; Robles et al., 2020). All these studies are single-country studies and do not provide insights in contextual variables influencing elite polarization and interaction, which might explain the differences in these findings.

Parliamentary relationships are influenced by the democratic model of the country. Lijphart (2012) describes two models of democracy. The majoritarian model is characterised by a legislature elected by a simple majority of the voters. The United Kingdom can be regarded as the majoritarian prototype, hence the alternative name "Westminster model". The second type of democracy, consensus democracy, usually employs proportional representation systems and leads to compromise and minority rights. Lijphart (2012) argues that the structures of power distribution represented by the consensus model fosters cooperation between politically dissimilar parties. Hence, consensus democracies are expected to exhibit a more densely connected parliamentary network. Conversely, other scholars postulate that political fragmentation is increased in proportional systems due to coalition forming and lower barriers of entry for smaller parties (Reynolds et al., 1999).

Comparative network analysis can provide insights in the influence of the democratic model on parliamentary interactions. Several authors argue that comparative network analysis presents a useful tool to address core questions in the social and political sciences (Vera and Schupp, 2006; Siegel, 2011; Fischer, 2011). Yet, while one-country Twitter studies are plentiful (see examples above), cross-country studies on parliamentary Twitter networks are sparse, with some notable exceptions. Urman (2020) emphasizes the importance of comparative research but focuses on mass polarization by means of audience duplication graphs. Van Vliet et al. (2020) introduce the Twitter Parliamentarian Database, including parliamentarians on Twitter in 26 countries, designed to foster comparative and transnational analysis. They developed a topology for retweets networks (Teernstra et al., 2018) and link this to the democratic system of a country (Van Vliet et al., 2020). Our study contributes to this existing work by applying a more systematic approach to compare network topologies and by integrating all layers of interaction on Twitter.

Twitter networks consist of three layers of interaction: the followers network, the retweets network and the mentions network. Each layer represents a different

type of communication. The followers network is a relational network, where an account is followed because of an interest in –and mostly, but not necessarily, agreement with– the account’s content. The followers network has shown to be very informative about ideological positions (Barberá, 2015). The retweets network is mostly a support network, resharing the tweets of users who think alike. Shi et al. (2017a) identify topical relevance, or congruence, as the most important factor in individual retweeting decisions. Several studies have found that party members are more likely to support or retweet candidates from their own party (Garcia et al., 2015; Cherepnalkoski and Mozetič, 2016; Esteve Del Valle and Borge Bravo, 2018; Van Vliet et al., 2020). In contrast, the mentions network is a more dialogical network that allows to interact with users who think differently. Parliamentarians have consistently be found to have cross-cutting interactions in the mentions network (Graham et al., 2016; Esteve Del Valle and Borge Bravo, 2018). This suggests that politicians are more likely to follow and retweet politicians with a similar ideology whereas they are more open to connect with opposing views in the mentions network.

### 4.3 Data collection

Our study includes 11 European countries with different political systems (Netherlands, Germany, United Kingdom (U.K.), Spain, France, Belgium, Italy, Romania, Poland, Ukraine, and Russia), and the United States (U.S.). Our choice to compare European countries to the U.S. is motivated by their dominant position in international politics and political research. For the aforementioned countries, all members of parliament (Chamber of Representatives and Senate), the president and members of cabinet (Prime minister, Ministers, Secretaries) and political parties (with seats in parliament as of May 2018) were collected from governmental websites and other internet sources, which are provided in Appendix 9.2.1. For each country, two independent coders with knowledge of the language and political context in the country were asked to manually check the Twitter handles of each politician, to select authentic accounts. The instructions that the coders received can be found in Appendix 9.2.1.1. This manual check was performed to avoid inclusion of incorrect (e.g. namesakes) or fake accounts (e.g. bots or identity impersonations (Goga et al., 2015)) in our dataset. Where the two coders did not agree on the correct Twitter handle, the Twitter handle of the politician was inspected by the authors. Using this list of Twitter handles, all tweets of the politicians’ accounts were streamed using the Twitter Stream API for the period of September 2018 till October 2019, resulting in one year of Twitter data.

An overview of the countries in our study can be found in Table 4.1. The Democracy Index (DI) for each country was derived from The Economist Intelligence Unit’s (EIU) Democracy Index 2019. The Democracy Index is based on five categories: electoral process and pluralism, civil liberties, the functioning of government, political participation, and political culture. The index lies between 0 and 10 and is based on the ratings for 60 indicators within these categories (EIU, 2020). The democratic model

of most countries is found in Lijphart (2012). Lijphart (2012) argues that democracies can be categorized among two dimensions. The *executives-parties* dimension groups five characteristics related to executive power, the party and electoral system, and interest groups. The *federal-unitary* dimension groups five characteristics related to federalism or unitary government. Based on these dimensions, consensus democracy is characterised by executive power sharing and decentralization, while majoritarian democracy is described by strong government and centralization. The electoral system is categorized by the International IDEA (2019) into three broad families: plurality/majority systems, proportional representation (PR) systems, and mixed systems. In a plurality/majority system, a candidate or party with a plurality of votes (i.e. more than any other) or a majority of votes (i.e. more than 50 percent) is elected. In a proportional representation system, the number of votes for a party correspond to the proportion of seats in an elected body. A mixed system combines a plurality/majoritarian voting system with an element of proportional representation (IDEA, 2019). Lastly, the party system was obtained from Bértoa (2020). In a multi-party system, multiple political parties form the government, whereas in a two-party system only two parties have a realistic chance of forming a majority, and in a dominant party system there is one party that has successively won the elections.

Table 4.1: Overview of the countries in our study.

Country	DI 2019 (EIU, 2020)	Democratic model (Lijphart, 2012)	Electoral system (IDEA, 2019)	Party system (Bértoa, 2020)
Netherlands	9.01	Consensual	Proportional	Multi
Germany	8.86	Consensual	Mixed	Multi
U.K.	8.85	Majoritarian	Plurality	Two
Spain	8.29	Majoritarian	Proportional	Multi
France	8.12	Majoritarian	Plurality (two rounds)	Multi
U.S.	7.96	Majoritarian	Plurality	Two
Belgium	7.64	Consensual	Proportional	Multi
Italy	7.52	Consensual	Mixed	Multi
Romania	6.49	Majoritarian (Gross, 2008)	Proportional	Multi
Poland	6.26	Consensual (Radecki, 2016)	Proportional	Two
Ukraine	5.90	Majoritarian (Christensen et al., 2005)	Mixed	Multi
Russia	3.11	Majoritarian (Chaisty, 2008)	Mixed	Dominant

#### 4.4 Methods

After some general insights on Twitter usage, activity and popularity for each of the countries, we will describe politicians' communication and relational networks on Twitter using social network analysis.<sup>1</sup> We will analyze followers, mentions and retweets networks separately, since they exhibit different properties with regard to the communication flow (Esteve Del Valle and Borge Bravo, 2018). We define a directed graph ( $G = (N, M)$ ) where the nodes ( $N$ ) represent politicians and the edges ( $M$ ) represent follower, mention or retweet relations on Twitter. A visualization of the

<sup>1</sup> Replication code can be found on Github: [https://github.com/SPraet/twitter\\_networks](https://github.com/SPraet/twitter_networks)

mentions and retweets network in all our countries at the party level can be found in 9.2.2. We first analyze the overall network structure of the parliamentary Twitter networks to measure “consensual unity” (Higley and Burton, 1989). Subsequently, we focus on inter-party communication as a measure of elite polarization.

#### 4.4.1 Network topology

To analyze the structure of social networks Himelboim et al. (2017) proposes a network-topology based on four network characteristics: density, modularity, centralization, and the fraction of isolated users. Unified networks are characterized by high density, low centralization, low modularity, and low fraction of isolates (Himelboim et al., 2017). For each country, we calculate these properties for the followers, mentions, and retweets network.

1. **Density.** This is the proportion of potential connections in a network that are actual connections and lies between zero and one (Jackson, 2010). This metric shows how connected politicians are in the network.

$$D = \frac{m}{n(n-1)} \quad (4.1)$$

With  $m$  the number of edges and  $n$  the number of nodes.

2. **Hierarchical structure/centralization.** Centrality can be measured using different approaches (see LD and Raj, 2017, for a comprehensive overview). We will use degree centrality, as it is widely used (Valente, 1996) and intuitive to understand. The degree centrality for a node  $v$  is the fraction of nodes it is connected to. The degree centrality of a network is defined as the sum of differences between the highest degree centrality and the degree centrality of all the other nodes in the network, divided by the maximum sum of differences (the latter can be proven to be equal to  $n - 3n + 2$ ) (Freeman, 1978).

$$C_D = \frac{\sum_i^n (C_D(v^*) - C_D(v_i))}{n - 3n + 2} \quad (4.2)$$

With  $C_D(v^*)$  the maximum degree centrality and  $C_D(v_i)$  the degree centrality of node  $i$ . This measure lies between 0 (very decentralized) and 1 (very centralized). In the case of a directed network, we can define two separate measures of degree centrality: inward hierarchy or outward hierarchy. Inward hierarchy is based on the in-degree (being followed, mentioned, or retweeted), while outward hierarchy is based on out-degree (following, mentioning, or retweeting other politicians).

3. **Modularity.** Modularity measures the strength of division of a network into different clusters or communities (Newman, 2010). Networks with high mod-



ularity have dense connections between the nodes within clusters but sparse connections between nodes in different clusters.

$$Q = \frac{1}{2m} \sum_{ij} (A_{ij} - \frac{k_i k_j}{2m}) \delta(c_i, c_j) \quad (4.3)$$

Where  $m$  is the number of edges,  $A$  is the adjacency matrix of  $G$ ,  $k_i$  is the degree of node  $i$  and  $\delta(c_i, c_j)$  is 1 if  $i$  and  $j$  are in the same community and 0 otherwise. To partition the graph in communities we make use of the Louvain algorithm (Blondel et al., 2008) which optimizes for modularity. The resulting modularity measures will lie between zero (the fraction of within-community edges is no different from what we would expect for a randomized network) and one (fully modular network).

4. **Isolates fraction.** Isolates are users who are not connected to other users in the network. In our case, these are users who have tweeted in the period under study, but did not mention/retweet others nor were mentioned/retweeted by others. The isolates fraction is the portion of isolates in the network and varies between 0 and 1.

$$I = \frac{n'}{n} \quad (4.4)$$

With  $n'$  the number of isolate nodes. The fraction of isolates allows to distinguish between two types of low-density networks: networks with small disconnected groups or networks with a high number of isolates.

After calculating these metrics for all countries, we apply hierarchical clustering analysis to learn which countries are more similar based on their Twitter network properties, and link this to the electoral and party system of the countries. First, we calculate the pairwise Euclidean distance between all countries, based on their network properties (i.e., the country's network values for density, inwards and outwards hierarchy, modularity, and isolates). We do this for the followers, mentions, and retweets networks separately. Next, we start an agglomerative clustering approach: each observation starts in its own cluster, and pairs of clusters are merged (based on minimum distance) in every step.<sup>2</sup>

---

<sup>2</sup> We use SciPy's hierarchical clustering <https://docs.scipy.org/doc/scipy/reference/generated/scipy.cluster.hierarchy.linkage.html>



#### 4.4.2 Inter-party communication

Secondly, we analyze inter-party communication as a measure of elite polarization along party lines. The External-Internal (E-I) index was developed as a measure of group embedding based on comparing the number of relations within groups and between groups (Krackhardt and Stern, 1988). It takes the number of connections (edges) of group members to outsiders, subtracts the number of connections to other group members, and divides by the total number of connections. In our case, politicians of the same party are considered as a group, and politicians from other parties are considered outsiders. The E-I index can be calculated as follows:

$$E - I = \frac{m_e - m_i}{m_e + m_i} \quad (4.5)$$

Where  $m_i$  denotes the number of internal connections (between two politicians from the same party) and  $m_e$  the number of external connections (between two politicians from a different party). The E-I index ranges from -1 (all connections are internal) to 1 (all connections are external). The proportion of external party relations is expected to be lower in the follower and retweets network than in the mentions network, because in the mentions network politicians more often interact with users with opposing views (Esteve Del Valle and Borge Bravo, 2018). Furthermore, it is to be expected that inter-party engagements are higher in consensual compared to majoritarian democracies.

Secondly, we investigate to what extent the relationships in parliamentary Twitter networks are in line with party ideology. We use the Left-Right Scale (RILE) by the Manifesto Project Dataset (Volkens et al., 2020) as an estimate of parties' left-right positions. The RILE index is a widely used method to measure left-right positions of parties. It measures how often a party references left (L) or right (R) issues in their electoral program (manifesto):<sup>3</sup>

$$RILE = R - L \quad (4.6)$$

The index lies between -100 (only left-wing issues) and +100 (only right-wing issues). The RILE scores for the parties in our study can be found in Appendix Table 9.8. Note that we do not have the scores for all parties available. Based on the RILE score, we calculate the Euclidean distance between parties in the two-dimensional ideological space. Next, we measure the number of inter-party relations for all pairs of political parties, divided by the total number of relations for each party. Finally, for each country, we calculate the Kendall rank correlation coefficient between the ideological distance and the proportion of inter-party relations between all pairs of

<sup>3</sup> see <https://manifesto-project.wzb.eu/down/tutorials/main-dataset>

parties. As an example, the ideological distance between German parties and the proportion of inter-party relations are shown in Table 4.2. We calculate the Kendall rank correlation using the vectorized matrices of ideological distance and inter-party relations, where we omit the diagonal elements (the distance to and relations within the own party). A scatterplot of both variables can be found in Figure 4.1. We expect party representatives to be more often connected to representatives of parties that are close in the ideological space than representatives of parties that are further away, especially in the retweets network. This corresponds to a negative rank correlation.

Table 4.2: Ideological distance (a) and proportion of inter-party follower relations (b) between parties in Germany.

(a) Ideological distance							(b) Proportion of inter-party follower relations						
	LINKE	SPD	90/Greens	FDP	CDU/CSU	AfD		LINKE	SPD	90/Greens	FDP	CDU/CSU	AfD
LINKE	0.00	20.48	20.86	42.49	44.67	59.34	LINKE	0.51	0.06	0.10	0.05	0.03	0.03
SPD	20.48	0.00	0.38	22.02	24.19	38.87	SPD	0.17	0.64	0.18	0.13	0.12	0.05
90/Greens	20.86	0.38	0.00	21.64	23.82	38.49	90/Greens	0.14	0.10	0.47	0.10	0.08	0.04
FDP	42.49	22.02	21.64	0.00	2.18	16.85	FDP	0.06	0.06	0.08	0.52	0.08	0.04
CDU/CSU	44.67	24.19	23.82	2.18	0.00	14.67	CDU/CSU	0.09	0.12	0.14	0.17	0.65	0.07
AfD	59.34	38.87	38.49	16.85	14.67	0.00	AfD	0.03	0.02	0.02	0.03	0.03	0.76




Figure 4.1: Scatterplot of the ideological distance and the proportion of inter-party relations between parties in Germany.

## 4.5 Results

### 4.5.1 Twitter usage

This Section provides some general insights on Twitter usage, activity and popularity for each of the countries. Table 4.4 and Figure 4.2 show the large country variation in degree of politicians with a (verified) Twitter account. In the US, almost all politicians have a verified Twitter account and to a lesser extent also in the Netherlands, Belgium, and France Twitter is popular amongst politicians. On the other side of the spectrum, in Ukraine, Romania, and Russia, Twitter is used by less than 30% of the politicians and almost none of the accounts are verified. Interestingly, this is not directly related to the popularity of the platform among the general public as Twitter is relatively well-used by the general population in Russia. Similarly, when looking at the average

number of tweets per month per politician, again, Ukraine, Romania and Russia are the least active (Figure 4.3).

Table 4.4: Overview of the number of politicians with a (verified) Twitter account per country

Country	Total number of politicians	Percentage on Twitter	Percentage verified	Average number of followers
U.S.	558	99%	94%	354,546
Netherlands	252	87%	38%	29,113
Belgium	383	87%	20%	9,844
France	962	83%	62%	24,765
Italy	971	80%	20%	24,994
Germany	794	73%	45%	16,039
Poland	577	70%	2%	16,782
Spain	630	56%	43%	49,017
U.K.	1486	53%	41%	46,143
Ukraine	450	34%	3%	41,505
Romania	487	28%	1%	2,064
Russia	654	24%	3%	54,755

More details on the total number of politicians are provided in Appendix Table 9.6

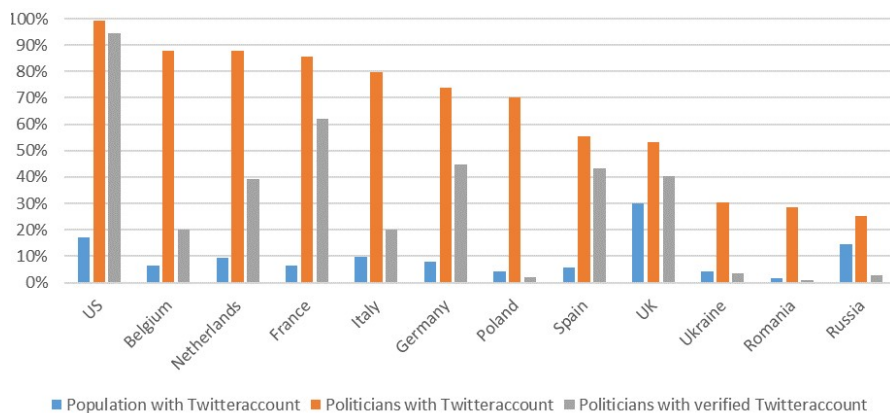


Figure 4.2: The proportion of politicians with a (verified) Twitter account compared to general Twitter use per country in 2019 (StatCounter, 2019).

Lastly, Figure 4.4 shows the average and median followers of the politicians' accounts as a percentage of the Twitter users per country. The follower counts are derived from the collected Twitter accounts using the Twitter API. On average, politicians have the highest (percentage) number of followers in the U.S. but this is mainly do to a few very popular accounts (e.g. in October 2019, @realDonaldTrump has 65,2 million followers and @POTUS has 26,8 million followers). Similarly, in Poland, some politicians are disproportionately popular on Twitter. On the other hand, the few politicians with a Twitter account in Russia and Ukraine are almost all followed by a relatively large number of users.




Figure 4.3: Average and median number of tweets in our dataset per politician per month.




Figure 4.4: Average and median number of followers of politicians' Twitter accounts, as a percentage of the Twitter users per country. Twitter users per country are derived by Twitter market share per country (StatCounter, 2019) and total population of the country (Worldbank, 2019).

#### 4.5.2 Network topology

Table 4.5, Table 4.6, and Table 4.7 show the network properties for the followers, mentions and the retweets networks respectively. As an example, the visualization of the three networks for Germany are shown in Figure 4.5. We will start with a comparison of these properties for the followers, mentions and retweets network. Next, we will discuss which countries are more similar based on their Twitter network properties based on the hierarchical clustering results.

As expected, the retweets network has higher **modularity** than the mentions network, but, more surprisingly, also higher modularity than the followers network, while follower and mentions network do not differ significantly.<sup>4</sup> Moreover, the detected clusters correspond better with the actual parties in the retweets network,<sup>5</sup> especially when modularity is high (Table 9.7). This indicates that politicians retweet mostly within their own parties, while party structure is less observable in the followers and mentions network. Mentions on the other hand foster interaction across the whole network and have therefore lower modularity.

Furthermore, the retweets network is the least **hierarchical**,<sup>6</sup> both in terms of retweeting and being retweeted, indicating the absence of dominant players. The higher inwards hierarchy in the mentions and followers network implies that certain ‘popular’ politicians are more frequently followed or mentioned than others, arguably a direct consequence of their official leadership position. Our findings are in line with Borge Bravo and Esteve Del Valle (2017), who found that having a central political position (e.g. party leader or government function) increases the centrality in the followers and mentions networks, but not so much in the retweets network. Centrality in the retweets network is more dependent on Twitter activity than official leadership. Lastly, we find that **density** is highest for the followers networks and lowest for the retweets networks. Politicians are thus most likely to follow and least likely to retweet fellow politicians. This finding is confirmed by the **isolates** fraction which is lowest for the followers networks and highest for the retweets networks.<sup>7</sup>

4 At the 5%-significance level after paired t-test with Bonferroni correction (Curtin and Schulz, 1998) for multiple comparisons ( $\alpha/3$ )

5 We use the Adjusted Mutual Information (AMI) score to measure how well the detected clusters correspond to the actual parties. The AMI returns a value of 1 when the two partitions are identical and 0 when their labels are independent. Using the paired t-test with Bonferroni correction again, we find a significant difference in AMI scores between mentions and retweets network, but not between follower and retweets network, nor follower and mentions network

6 Using the paired t-test with Bonferroni correction, we find a significant difference between mentions and retweets network, and between follower and retweets network, but not between follower and mentions network

7 Again, at the 5%-significance level after paired t-test with Bonferroni correction

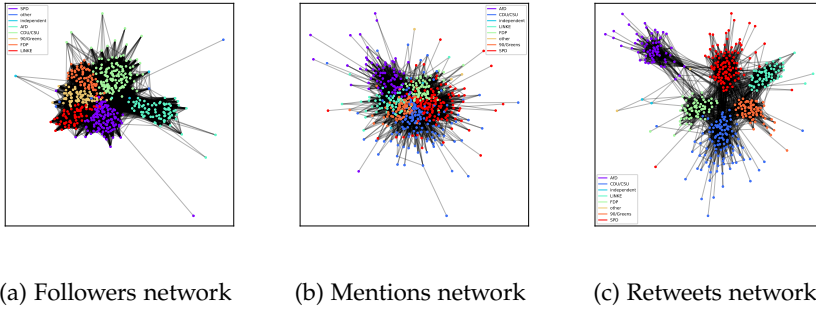


Figure 4.5: The parliamentary followers (a), mentions (b), and retweets (c) network in Germany.

Table 4.5: Description of the followers networks of the 12 countries in our study.

Country	Density	Hierarchy in	Hierarchy out	Modularity	Isolates
Netherlands	0.24	0.42	0.58	0.13	0.03
Germany	0.14	0.46	0.66	0.40	0.05
U.K.	0.11	0.22	0.58	0.36	0.07
Spain	0.14	0.41	0.29	0.47	0.03
France	0.19	0.56	0.74	0.26	0.02
U.S.	0.15	0.30	0.60	0.32	0.08
Belgium	0.18	0.47	0.33	0.31	0.02
Italy	0.05	0.25	0.38	0.37	0.03
Romania	0.02	0.16	0.11	0.46	0.32
Poland	0.19	0.54	0.57	0.32	0.02
Ukraine	0.06	0.45	0.19	0.19	0.13
Russia	0.06	0.62	0.45	0.14	0.17

Table 4.6: Description of the mentions networks of the 12 countries in our study.

Country	Density	Hierarchy in	Hierarchy out	Modularity	Isolates
Netherlands	0.10	0.40	0.32	0.20	0.09
Germany	0.04	0.43	0.22	0.35	0.08
U.K.	0.04	0.49	0.23	0.23	0.11
Spain	0.05	0.62	0.26	0.40	0.06
France	0.05	0.63	0.18	0.22	0.03
U.S.	0.02	0.50	0.14	0.31	0.14
Belgium	0.08	0.57	0.34	0.28	0.05
Italy	0.01	0.26	0.08	0.35	0.39
Romania	0.00	0.04	0.06	0.67	0.84
Poland	0.07	0.63	0.30	0.21	0.13
Ukraine	0.01	0.12	0.05	0.47	0.74
Russia	0.00	0.06	0.04	0.75	0.75

Table 4.7: Description of the retweets networks of the 12 countries in our study.

Country	Density	Hierarchy in	Hierarchy out	Modularity	Isolates
Netherlands	0.04	0.12	0.10	0.61	0.19
Germany	0.02	0.11	0.12	0.71	0.18
U.K.	0.03	0.19	0.20	0.44	0.17
Spain	0.04	0.23	0.17	0.69	0.12
France	0.03	0.39	0.23	0.28	0.10
U.S.	0.01	0.13	0.15	0.44	0.25
Belgium	0.03	0.16	0.23	0.66	0.11
Italy	0.00	0.12	0.10	0.56	0.46
Romania	0.00	0.02	0.03	0.58	0.91
Poland	0.04	0.32	0.19	0.49	0.21
Ukraine	0.00	0.07	0.05	0.52	0.80
Russia	0.00	0.18	0.04	0.46	0.72

Next, we continue with a comparison of the network properties between countries. The dendrogram in Figure 4.6 displays the results of the hierarchical clustering based on the followers, mentions, and retweets network topologies of the 12 countries, as well as for the three network layers combined. The y-axis shows the Euclidean distance between clusters of countries. With respect to **follower** relations (Figure 4.6a), Romania is a clear outlier. As we learned from Figure 4.4, Romanian politicians do not have many followers. The isolates fraction for the followers network is very high and density is low. For the mentions and retweets network, Russia, Romania and Ukraine —also the least democratic countries— are clustered separate from the other countries. Twitter activity is low in these countries (see Figure 4.3) and thus density of the mentions and retweets networks is extremely low and there are many isolates. The network topologies of the U.K. and the U.S. —both two-party systems— are very similar to each other for all three network settings. In the followers network, they have a relatively low inwards hierarchy compared to the multi-party countries. Possibly because in a multi-party system the representatives of larger parties are more frequently followed than representatives from smaller parties. Only Italy appears to be closer to the two-party system in that respect. The Netherlands, France, Germany and Poland all have high density and high in- and outwards hierarchy in the followers network, while Belgium and Spain have a lower outwards hierarchy.

In the **mentions** network (Figure 4.6b), Germany and the Netherlands have, similar to the two-party systems, low inwards hierarchy. Mentions are more equally distributed among all parliamentarians than in the other countries. The other multi-party systems have higher inward hierarchy, which means there are some central accounts that are more often mentioned compared to the large amount of parliamentarians with little political influence. These central accounts are likely leading cabinet positions or party leaders (of larger parties).<sup>8</sup> Furthermore, consensual democracies have higher density in the mentions network than majoritarian systems.<sup>9</sup> This is in line with Lijphart (2012) who suggests that consensus democracies lead to increased cooperation and dialogue. Spain leans slightly over to a majoritarian system but does have proportional representation. The exceptions to this rule would be France —majoritarian system with high density—, and Germany —consensus democracy with low density. Also, Italian politicians do not mention others frequently and thus the network has a low density. Alternatively, another possible explanation for high density in the network is a relatively low number of politicians in parliament, making it ‘easier’ to have a densely connected network (again, France being the exception to the rule).

Regarding **retweets**, networking behavior and the resulting clusters are slightly different (Figure 4.6c). As mentioned before, retweets reflect endorsement and thus the patterns of retweeting can be revealing of the political alliances within a

<sup>8</sup> For example, in Spain the most often mentioned accounts are —next to party accounts— Pablo Casado Blanco (party leader PP, largest opposition party), Jaime de Olano (Deputy Secretary General PP), Albert Rivera (Party leader Ciudadanos), and Teodoro García Egea (Secretary-General PP)

<sup>9</sup> Tested with two-sample t-test



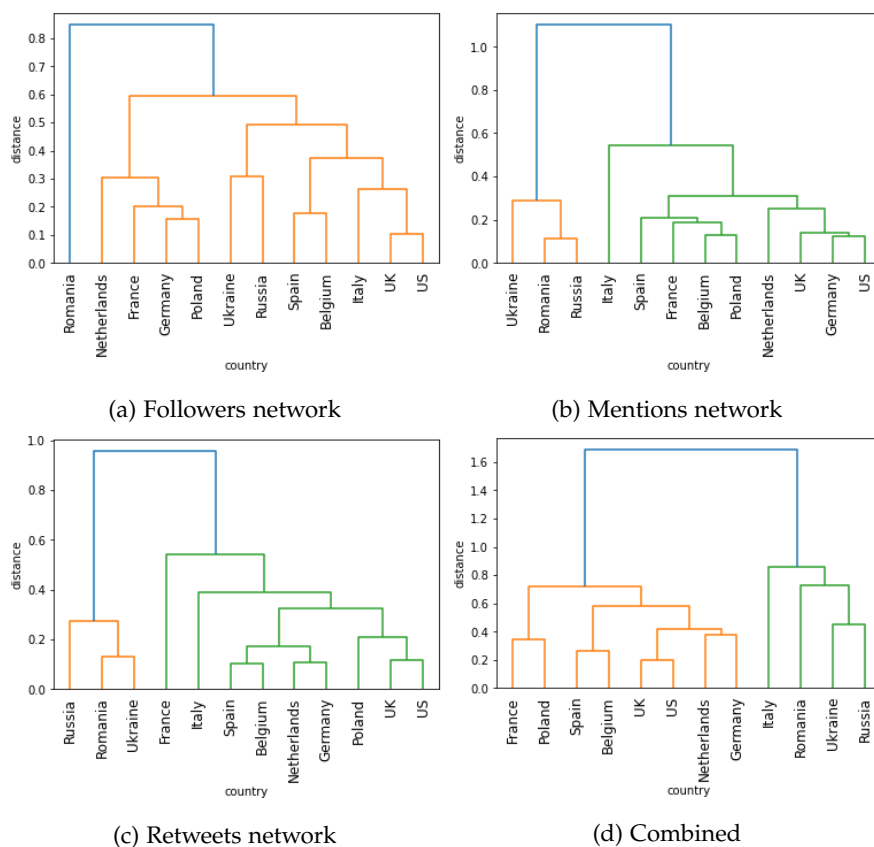


Figure 4.6: Hierarchical clustering dendrogram of the topology of the parliamentary followers (a), mentions (b), and retweets (c) networks, as well as a combination of all three (d) for the 12 countries in our study.

country. Systems with proportional representation have higher modularity (with an exception for Poland). This finding is in accordance with scholars arguing that political fragmentation is increased in proportional systems due to coalition forming and lower barriers of entry for smaller parties (Reynolds et al., 1999). Also Italy is highly modular, but more distant from the other proportional systems because of low density. France has by far the lowest modularity which results in a dense and interconnected retweets network. The U.S. and U.K. show a strong two-party structure in the retweets network, characterised by low modularity and high in- and outwards hierarchy. Likewise, in Poland the network is dominated by the two major parties (see Figure 4.7), reflecting the recent polarization of the party landscape into two competing “blocks” of parties (Tworzecki, 2019). The groups we can distinguish using this quantitative hierarchical clustering approach on the retweets network largely correspond to the network archetypes of Van Vliet et al.


(2020), that were derived qualitatively. The groups are visualized in Figure 4.7.<sup>10</sup> We find a cluster of countries with bipolar networks (U.S., U.K., and Poland), fragmented networks (Germany, Netherlands, Spain, Belgium, and Italy), and a cohesive network (France). However, based on our clustering results, we do not differentiate between what Van Vliet et al. (2020) call “networks with rogue clusters” and “fragmented networks”. A contribution of this study is that three less-democratic countries were also included which resulted in an additional archetype: “unconnected networks”, characterized by low density and many isolates.<sup>11</sup>

Combining the information of all three networks results in an almost perfect representation of the democratic systems and functioning of the countries. Russia, Romania, and Ukraine are clustered together. Politicians of these countries interact the least with other politicians on Twitter, and these countries have the lowest democracy scores. Yet, Figure 4.2 and Figure 4.4 show that Twitter is actually used in these countries and that politicians are followed by the population. Poland is an exception, with a low democracy score but a lot of political interaction on Twitter, and Italy is a more democratic country with little activity in the mentions and retweets network. The U.K. and the U.S. are plurality two-party systems and have a more equal distribution of mentions and less modular retweets network. Germany and the Netherlands, the most democratic countries, are proportional systems but also have low inwards hierarchy in the mentions network. Moving from high to low democracy score, Spain and Belgium are the next proportional systems. They have high modularity in the retweets network but also fairly high density and low isolates in the mention and retweets network. Finally, Poland is a proportional system that shows a two-party structure with relatively low modularity, while France is a plurality system with relatively high density.


<sup>10</sup> After excluding isolates, we used NetworkX Spring Layout ([https://networkx.org/documentation/stable/reference/generated/networkx.drawing.layout.spring\\_layout.html](https://networkx.org/documentation/stable/reference/generated/networkx.drawing.layout.spring_layout.html)) to visualize the network. The positions of the nodes are optimized using Fruchterman-Reingold force-directed algorithm (Fruchterman and Reingold, 1991). The algorithm finds an equilibrium between two opposing forces: edges hold nodes close, while nodes repel other nodes. This way, connected nodes are positioned closer together in the visualization than unconnected nodes.

<sup>11</sup> Since we exclude isolates from the network visualization, the resulting visualizations consists of very few nodes.


### Bipolar networks



(a) UK




(b) US




(c) Poland


### Fragmented networks




(d) Germany




(e) Netherlands



(f) Spain




(g) Belgium



(h) Italy


Figure 4.7: Network visualization for the retweets networks of the 12 countries in our study.

**Cohesive network**




(i) France


**Unconnected networks**



(j) Russia



(k) Romania



(l) Ukraine

Figure 4.7: (Continued) Network visualization for the retweets networks of the 12 countries in our study.

4.5.3 *Inter-party communication*

The E-I index measures the relative amount of external party communication and ranges from -1 (all ties are internal to the party) to 1 (all ties are external to the party). Table 4.8 shows the average E-I index per country. For most countries, the E-I index is lowest for the retweets network.<sup>12</sup> For the follower and mentions network, external party relations are correlated with the effective number of parties. A higher number of effective parties requires more cooperation between different parties (e.g. coalitions), which results in parliamentarians following and mentioning—but not retweeting—candidates from other parties more frequently. Interestingly, Belgium and the Netherlands even have a positive E-I index, which means they follow or mention politicians from other parties more often than from their own party.

Table 4.8: E-I index per country, ranked from low to high effective number of parties (Bértoa, 2020).

Country	Effective number of parties	E-I Followers	E-I Mentions	E-I Retweets
Russia	1.7	-0.85	-1.00	-0.67
U.S.	2.0	-0.66	-0.45	-0.54
U.K.	2.5	-0.60	-0.27	-0.80
Poland	2.8	-0.27	-0.06	-0.50
France	3.0	-0.28	-0.13	-0.55
Romania	3.5	0.29	-0.24	-0.64
Spain	4.2	-0.50	-0.23	-0.17
Italy	4.3	-0.00	-0.01	-0.96
Ukraine	5.5	0.36	-0.15	-0.82
Germany	5.6	-0.33	-0.09	-0.73
Belgium	7.8	0.11	0.22	-0.18
Netherlands	8.1	0.44	0.37	-0.87
Kendall's tau		0.64	0.70	0.36
P-value		0.00	0.00	0.12

Next, we investigate whether parliamentarians prefer to interact with parties that are ideologically close to their own party. Our results for the Kendall rank correlation between ideological distance and proportion of inter-party relations are inconclusive (Table 4.9). Only the correlations for the Netherlands are significant at the 0.05 level for all networks.<sup>13</sup> For all three Dutch networks, the further apart the parties are ideologically, the less they will interact on Twitter. Additionally, for the followers network in Germany and Romania, and for the retweets network in Belgium we

<sup>12</sup> Paired t-test with Bonferroni correction

<sup>13</sup> After Bonferroni's correction for multiple (three) comparisons

find similar results. For the other countries, the number of parties<sup>14</sup> is too low to find significant results. For Ukraine and US only the RILE scores of two parties are available, hence the correlation cannot be calculated. Romania and Russia have nearly no interactions in the mentions and retweets network. Additional research is necessary to reveal which cross-party interactions most frequently take place and for which purpose. Do politicians use Twitter as an instrument to challenge and criticize the opponent or do they rather interact with ideologically similar parties? A more in-depth sentiment analysis of the tweets could provide more clarity on this.

Table 4.9: The Kendall rank correlation coefficient between inter-party relations and ideological distance (RILE) (Volkens et al., 2020).

Country	Followers	Mentions	Retweets
Belgium	-0.20	-0.10	-0.39***
France	0.04	0.10	0.00
Germany	-0.34**	-0.12	-0.27
Italy	0.02	0.13	0.08
Netherlands	-0.19***	-0.18***	-0.26***
Poland	0.15	0.19	-0.02
Romania	-0.37**		
Russia	-0.18		
Spain	-0.06	0.06	0.14
U.K.	-0.06	0.12	-0.08
Ukraine			
U.S.			

\*p < .1/3; \*\*p < .05/3; \*\*\*p < .01/3  
Kendall's tau.

## 4.6 Conclusion

Elite polarization and the amount of cooperation among the elite have important implications for our democracy (Druckman et al., 2013; Higley and Burton, 1989). The rise of social media has altered existing political relations and simultaneously offered new opportunities to empirically analyze these structures. A plethora of studies explore political polarization on social media, with sometimes contradictory results. A possible explanation for these contradictions is the (institutional) context in which the study takes place. Yet, little research has focused on structured comparison across multiple countries to gain insights in contextual variables and country characteristics.

We analyze the interactions in 12 parliamentary Twitter networks and find that the network topology is related to the democratic functioning and political system of

<sup>14</sup> The number of parties for which we have the RILE score, i.e. the number of parties that overlap in our study and that of Volkens et al. (2020)

the countries in our study. Consensual democracies are characterised by more dense relations but also higher hierarchy and higher fragmentation in the retweets network, while systems with plurality voting generally have lower modularity. Parliaments with a high effective number of parties are more cooperative, which results in higher inter-party relations. By design, two-party systems exhibit higher elite polarization on Twitter. In fact, these findings are far from novel or unexpected, and correspond to established theoretical insights in the field (such as Lijphart, 2012). However, the prominent empirical confirmation of these theoretical concepts highlights the importance of including institutional context in social media research. We need more comparative research to truly understand the influences on and the effects of polarization in our society.

Secondly, we show differences in the followers, retweets and mentions networks that hold across all countries and political systems. The retweets network is most polarized or fragmented, while politicians engage more often in inter-party interactions in the followers and mentions network, especially in countries with high effective number of parties. Twitter can be conducive to both cross-cutting interactions and echo chambers depending on the layer of interaction. Furthermore, not all interactions are necessarily positive for democracy as Twitter can be used to permanently follow and attack the communication of a political opponent. Again, this could be an important part of the explanation why we find contradictory results on the effect of social media on polarization. The type of interactions we undertake on social media determines its polarizing effect.

In this work, we have specifically focused on a network approach. Nonetheless, we do want to emphasize that social network analysis in combination with textual analysis can provide more detailed insights in the motivations or goals behind interactions. For example, sentiment analysis can uncover whether a mention is meant to criticize or support an opponent (Khatua et al., 2020), and with topical analysis we can learn how politicians communicate about certain issues, and which topics induce controversy (Al-Ayyoub et al., 2018). Furthermore, the scope of this study is limited to interactions among parliamentarians. Nonetheless, Twitter is used by politicians to communicate not only with other politicians but also with citizens, opinion leaders, and journalists (Jungherr, 2016; Vargo et al., 2014). Therefore, including the interactions with other Twitter users could provide additional insights, and could be relevant for future research.

Our results indicate that both the institutional context as well as the platform layer should be taken into account when trying to understand parliamentary interactions and elite polarization. Given the effects of elite polarization on mass polarization and the importance of elite cooperation for the democratic functioning of a country, these findings can have far-reaching consequences to improve our understanding of these phenomena. We show how social network analysis could be a fruitful opportunity for future comparative research on politics and social media.





## Part II

### FACEBOOK LIKES TO STUDY LIFESTYLE AND POLARIZATION



## Facebook Likes to study lifestyle politics

---

“Lifestyle politics” suggests that political and ideological opinions are strongly connected to our consumption choices, music and food taste, cultural preferences, and other aspects of our daily lives. With the growing political polarization this idea has become all the more relevant to a wide range of social scientists. Empirical research in this domain, however, is confronted with an impractical challenge; this type of detailed information on people’s lifestyle is very difficult to operationalize, and extremely time consuming and costly to query in a survey. A potential valuable alternative data source to capture these values and lifestyle choices is social media data. In this study, we explore the value of Facebook Like data to complement traditional survey data to study lifestyle politics. We collect a unique dataset of Facebook Likes and survey data of more than 6,500 participants in Belgium, a fragmented multi-party system. Based on both types of data, we infer the political and ideological preference of our respondents. The results indicate that non-political Facebook Likes are indicative of political preference and are useful to describe voters in terms of common interests, cultural preferences, and lifestyle features. This shows that social media data can be a valuable complement to traditional survey data to study lifestyle politics.

## 5.1 Introduction

Voting is central to the democratic process and the well-functioning of our political system. Therefore, voting behavior and party preference are well-studied by a broad range of social scientists who try to explain how and why decisions are made by the electorate. Traditionally, to understand individuals' votes, scholars relied on socio-structural factors such as group identification, religious affiliation, and socio-economic status. Electoral pioneers from the Michigan and Columbia school both stressed the importance of these long-term factors in explaining voting behavior (Berelson et al., 1954; Campbell et al., 1960). In the U.S. context party identity was key, while in many European countries, more religious and class-based cleavages were driving voters (Lipset and Rokkan, 1967). As these stable factors started losing importance, scholars gradually started devoting more attention to short-term factors such as concrete issues and popular candidates (Dalton et al., 2000; Dalton and Wattenberg, 2002). Studies showed that a better educated electorate became more volatile and more affected by issue priorities (Petrocik, 1989) and candidate evaluations (Funk, 1999).

"Lifestyle politics" represents an alternative way to understand political preference and voting behavior. Authors such as Bennett (1998) and Giddens (1991, 2013) argue that in our post-modern society, personal identity is replacing collective identity and that individuals increasingly let their personal politics depend on lifestyle choices. Politics focuses more and more on broad issues like identity, values, and moral and social orientation (Giddens, 2013). According to DellaPosta et al. (2015, p. 1474) we "are increasingly likely to find our local communities and social networks populated by individuals with similar aesthetic tastes, leisure activities, consumer preferences, moral practices, and ways of life." Our attitudes on climate change, religion, migration, sexual minorities or workers' welfare are strongly connected to our consumption choices and daily lifestyle (Purhonen and Heikkilä, 2017); in turn, lifestyle is becoming more important in understanding individuals' political opinions and voting behaviors. In recent studies, political preference has been associated with leisure activities and personal tastes (DellaPosta et al., 2015), brand choice in supermarkets (Khan et al., 2013), sustainable behaviors (Kidwell et al., 2013), food and music taste (Purhonen and Heikkilä, 2017; Shi et al., 2017b) and movie preferences (Roos and Shachar, 2014). Although the linkage between politics and many aspects of our personal life were documented, what is driving this alignment between ideological preferences and seemingly unrelated lifestyle dimensions is less clear. While some suggest that these lifestyle choices are a way to express one's (political) identity (Roos and Shachar, 2014), others rather stress the deeper role of social and cultural homogeneity (DellaPosta et al., 2015). Furthermore, we know relatively little on whether these "cultural fault lines" that are clearly visible in the two party system in the US (Shi et al., 2017b) are as present and determining in a multi-party system with a much broader variety of ideological players.

The growing political polarization renewed the issue of “lifestyle enclaves” as a prime interest for a wide range of social scientists (Iyengar et al., 2012; Lelkes, 2016). Yet, detailed information on individuals’ lifestyles is very difficult to collect, which complicates empirical and comparative studies in this domain. The options of leisure activities, movies, music, cultural activities, etc. that could be included in survey questions on lifestyle politics are endless, therefore, including these type of questions will not only occupy a lot of survey space and time but will also result in non-exhaustive option lists. Moreover, it implies that researchers are able to deductively make a selection of the lifestyle indicators (i.e. cultural places, movies, products) that should be included to best explain or predict vote choice.

A potential valuable alternative data source to capture these cultural and lifestyle practices is social media data. Today, thanks to the Internet and social media, an unseen amount and granularity of data are available. People visiting webpages or liking Facebook content leave little “bread crumbs” behind in the digital world that are indicative of their interests and personality. From this fine-grained behavioral data, inferring unknown information about a user is possible by applying predictive modeling techniques (Martens et al., 2016). In this study, we explore the potential of Facebook Like data for capturing lifestyle and predicting political and ideological preference. With Facebook Likes we refer to the mechanism used by Facebook users to express their positive association with public Facebook pages of products, sport clubs, musicians, books, restaurants, etc.

Kosinski et al. (2013) showed that Facebook Like data can be used for predicting personality traits and political attitudes. Since then, researchers affirmed the potential of Facebook data for predicting individual political orientation (David et al., 2016; Kristensen et al., 2017; Chiu and Hsu, 2018; Bach et al., 2019). Some scholars combined these new social media data with traditional survey data to study and understand political behavior. For example, Bond and Messing (2015) estimate the ideology of politicians and their supporters using individual citizens’ Facebook Likes of political figures to study the relationship between ideology and age, social relationships and ideology, and the degree of polarization among the electorate. Eady et al. (2019) apply a method developed by Barberá (2015), to quantify the ideological distributions of users’ online political and media environments on Twitter and study the extent to which liberals and conservatives live in so-called “echo-chambers”. Some other examples include comparing individual exposure to news and politics content (Wells and Thorson, 2017) and understanding news sharing behavior (Joseph and Wihbey, 2019; Mosleh et al., 2020).

Although the combination of survey data and digital trace data seems to offer valuable insights in political attitudes and behavior, relatively few studies actually combined both types of data sets (Stier et al., 2019). In particular, our inclusion of non-political Facebook pages to study lifestyle politics is innovative, as most previous literature mainly worked with the inference of ideology based on the accounts or pages from political actors or news media.

To explore the value of digital trace data to complement classical survey research, we gathered Facebook Likes and survey data of more than 6,500 participants in Flanders, the Dutch-speaking part of Belgium. The data gathering was of course done with user consent and a clear privacy statement. Based on these Facebook Likes, we built models to predict political and ideological preference, and compare this to predictive models based on traditional survey data. The interpretation of these models can be used to gain insights into voter profiles. The main contributions of this study are twofold. First, we explore the use of Facebook Like data to complement and improve traditional survey data to study lifestyle politics, with an explicit focus on the contribution of non-political likes. Second, we collect a unique dataset of Facebook Likes and survey data to gain insights into voter profiles in Belgium, a fragmented multi-party system.

## 5.2 Predictive modeling with Facebook Likes

Facebook Likes<sup>1</sup> express a positive attitude or interest, and are comparable to webpage visits, purchase behavior, payments or location data. For example, observing a users' Facebook Likes related to books provides similar information to watching someone's book closet or a list of purchased books online. Liking a Facebook page signals a user's desire to see more posts from the page's publisher. Additionally, a Facebook Like has been conceptualized as a form of social endorsement, since these likes are publicly observable by friends in one's network (Bond and Messing, 2015). Since like behavior is observable, it is used as a form of self-expression, and in line with the theory of Goffman (1967) it is believed that users are building an idealized version of themselves. At the same time, this "ideal" self should remain congruent with how one is perceived in the offline world, as Facebook networks are often grounded in offline relationships (Eranti and Lonkila, 2015). Indeed, in a study on consumer identity, Hollenbeck and Kaikati (2012) find that users like Facebook pages (of brands) to present –an ideal version of– themselves on the platform. Facebook Likes have been proven powerful to infer interests and psychological traits, and became invaluable for user profiling and personalized advertising applications (Matz and Netzer, 2017). More concretely, Piazza et al. (2017) investigated the relationship between Facebook Likes and individual lifestyle, and they found a correlation between the activity, interests, and opinions of an individual and their Like information.

We argue that Facebook Likes could be a valuable addition to survey data for electoral research because of at least four different reasons: (1) through the direct measurement of actual behavior we avoid recall error, and subjectivity<sup>2</sup> that are

<sup>1</sup> We refer to *Public Page Likes*, i.e. the public Facebook pages that a user likes and that show up as being liked in the About section of that person's profile (see [https://www.facebook.com/help/171378103323792?helpref=uf\\_permalink](https://www.facebook.com/help/171378103323792?helpref=uf_permalink)). We do not include likes or emotional reactions (Love, Haha, Wow, Sad, and Angry) to Facebook posts.

<sup>2</sup> For example, Guess et al. (2019a) linked original survey data with respondents' observed social media data to validate self-reports of political activity; and discovered a substantial discrepancy between objective

specific to answering survey questions (Furnham, 1986), (2) it requires less effort from the respondents, and thus it is less time-consuming and costly than collecting survey data, (3) they can provide an unfiltered look and unique information on interests and lifestyle that we cannot grasp (fully) with survey-based data and (4) they can inductively help to identify the most important indicators of a phenomenon (such as lifestyle politics in the case of our study), which can then later be included in a survey.

As discussed in Chapter 1, we must keep in mind that online data is a proxy for the real behavior under study, and thus, drawing inference and generalizing results from this type of data must be performed with caution (Nagler and Tucker, 2015; Dalton, 2016). For example, what you like on Facebook is not necessarily what you like in real-life and thus true behavior can only be studied indirectly (Dalton, 2016). More specifically, two potential hidden biases need to be considered. First, Facebook Like behavior may be affected by user induced biases such as social desirability and intentional misrepresentation (creating the “ideal” self). Second, Facebook profile data are affected by the mechanics of the platform,<sup>3</sup> such as the personalization by Facebook algorithms (Kosinski et al., 2015). Furthermore, behavioral data can suffer from exhibiting a low signal-to-noise ratio, since the behavior we capture can be unrelated to the target question. Surveys, if properly designed and implemented, provide better quality controls and allow for more targeted questions and responses (Buntain et al., 2016). A last issue with behavioral data is privacy. Since these data are often very personal and sensitive, the data must be collected and stored with respect for users’ privacy, addressing challenges on user consent, data anonymization, secure storage, etc. (Zimmer, 2010).<sup>4</sup> Yet, because of the enormous amount of information available in behavioral data, there is much to learn from it when following a rigorous research approach (Nagler and Tucker, 2015; Dalton, 2016).

We will use a predictive modeling approach to study Facebook Likes and political preference. Although most empirical political science research relies on explanatory modeling to test theory-driven explanations (Druckman et al., 2006), predictive modeling is more suitable to uncover complex patterns from data that might lead to the generation of new hypotheses (Cranmer and Desmarais, 2017). The distinction between explanatory and predictive modeling exhibits some practical implications to each step of the modeling process (Shmueli et al., 2010), as explained in Chapter 1.

### 5.3 Data collection

The data collection started in March 2018 and focused on Flanders, the Dutch-speaking part of Belgium, representing around 60% of the population. A detailed

---

and self-reported posting behavior, which could be due to a subjective interpretation of what is considered to be “political”.

<sup>3</sup> The “social media logic” we refer to in Chapter 1

<sup>4</sup> We elaborate on our research design with regard to these issues in Chapter 2

survey with questions on socio-demographics, media consumption, political preference and attitudes was sent to 4,500 respondents online. Surveying was done by Dynata/SSI, who distributed the survey among their own online panel, which consists of a diverse collection of citizens in terms of age, gender, and educational level.<sup>5</sup> Targets were set on these socio-demographic characteristics to increase the representativeness of the sample.<sup>6</sup> Of these respondents, 524 agreed to provide us access to their Facebook Like data via Facebook Login. In May–June 2018, a second round of data collection was conducted, where we disseminated a shorter survey and Facebook Login through the online webpages of popular Flemish newspapers. We asked people to use our tool, which would predict their ideological position based on their personal Facebook Likes. An additional 6,209 respondents agreed to provide access to their Facebook data, and they completed 12 survey questions about their media consumption and political preference. The Facebook Likes and survey questions were collected with user consent and stored anonymously and securely on a local server. The data are used for scientific purposes only and will under no circumstances be shared with other institutions or companies. Results will be shown on an aggregate level only and participants have the right to stop their collaboration at any time and ask for their data to be removed. More details about the data collection and a full discussion of the privacy and ethical concerns can be found in Chapter 2.

Per user, we stored the name and timestamp of all public Facebook pages they liked. This resulted in a total of 595,994 unique Facebook pages. For privacy reasons (see Appendix 1.3.4), only pages that are liked by a minimum of 30 respondents in our dataset will be reported later in the analysis, which results in 10,226 pages. For these pages we searched for the category that was assigned to them on Facebook. In total, Facebook shows over 1,300 different page categories.<sup>7</sup> We started from the categorization that Facebook uses to classify public pages, but we refined and adjusted it into 20 categories (see Table 5.1). Note that our categories are not mutually exclusive; for example, an art festival can be included in the category **Arts & Culture** as well as in the category **Event & Festival**.

The survey data (see Appendix 9.3.1.2) include some of the basic variables that are generally included in models of voting behavior in Belgium (see for instance Delwit et al., 2015; Hooghe and Dassonneville, 2018). Specifically, we added several of the most important structural determinants of ideology and voting behavior (gender, age and education), as well as the use of different media sources (TV, newspaper, social media, etc.), interest in the news (sports, culture, home affairs, foreign affairs, etc.), and general interest in politics.

<sup>5</sup> For more information on Dynata and their panel see <https://www.dynata.com/>.

<sup>6</sup> With regard to age our sample of 4500 respondents is slightly older than the population.

<sup>7</sup> An overview of all possible Facebook categories can be found on <https://www.facebook.com/pages/category/>.



Table 5.1: Description and number of pages for the 20 Facebook categories.

Category	Description	# pages
Communities	Communities, interests and places	1,610
Companies & Business	Companies, entrepreneurs, stores, shops, etc.	1,468
Music	Music, bands, producers, record labels, albums, etc.	1,454
Apps, Websites & Blogs	Apps, websites and blogs	979
Products & Services	Products, brands, financial services, marketing, etc.	912
Artists & Public Figures	Artists and public figures	743
News & Media	News, media, radio, magazines, etc.	669
Games, Humor & Entertainment	Games, humor, amusement, comedy, etc.	615
Tv Shows	TV shows and episodes	528
Sports & Health	Sports, athletes, gym, health	488
Civil Society	Nonprofit organizations, labor unions and religious organizations	483
Politics	Politicians, political parties and government organizations	468
Arts & Culture	Arts, culture, photography, museums, etc.	386
Food, Drinks & Restaurants	Food, cooking, restaurants, breweries, etc.	366
Movies	Movies, films, actors and cinema	352
Events & Festivals	Events, festivals and concerts	291
Bars, Cafes & Night clubs	Bars, cafes, pubs, clubs etc.	224
School, University & Education	Schools, universities, student organizations and education	195
Books & Authors	Books, libraries, publishers, writers	193
Travel	Travel, tour agencies and tourism	138
TOTAL		10,226

Two survey questions will be used as target variables. The number of participants per target variable can be found in Table 9.10 in the Appendix.

1. *Ideological leaning.* The participants positioned themselves on a scale of 0 (most left) to 10 (most right). In the analysis we consider the numbers 0 to 3 as “left”, 4, 5 and 6 as “center”, and numbers 7 and higher as “right”.
2. *Party preference.* The participants indicated how likely they are to ever vote for each of the seven main Flemish political parties on a scale of 1 (never vote for party) to 10 (definitely vote for party): the worker’s party (PVDA), the green party (Groen), the social democratic party (Sp.a), the Christian democratic party (CD&V), the liberal party (Open VLD), the Flemish nationalist party (N-VA), and the extreme right party (Vlaams Belang, VB). The preferred party per participant is the one with the highest score. In case of a tie between two or more parties, all tie-parties are considered equally important.

Our sample of Facebook users consists of a diverse mix of users in terms of gender, age and education levels. Yet, our sample contains less females, less participants older than 55 and less lower educated participants compared to the general population. Weights were applied to our survey samples (see Appendix 9.3.1.3) but this did not influence the results,<sup>8</sup> therefore we will only report the unweighted results hereafter. Moreover, through the self-selection of participants, Facebook users with a higher political interest are overrepresented. Likewise, some parties are over- or underrepresented in our sample. However, since the goal of our study is not to predict aggregated election results but rather to gain insights into voter groups, this is not a particular stumbling block in this study. Nevertheless, it is likely that we will achieve more accurate results for political parties that are well-represented in our sample.

## 5.4 Methods

As discussed in Section 5.2, our methodology consists of predictive modeling to study lifestyle politics. We apply predictive models to infer the political and ideological preference of our respondents, based on their Public Page Likes on Facebook.<sup>9</sup> Facebook pages are encoded as dummy variables where the value 1 indicates the user has liked the page and 0 indicates the user has not liked the page. The goal of these models is to optimize prediction accuracy at the individual level. Yet, the coefficients of these models reveal which likes are predictive for a certain political preference, providing insight in the interconnection between lifestyle and politics. To show Facebook Likes are not just capturing socio-demographics in an indirect way, we compare the models built on Facebook Likes to models built on survey data

<sup>8</sup> The spearman rank correlation between the coefficients of the models with and without survey weights applied was 0.94

<sup>9</sup> Replication code can be found on Github: [https://github.com/SPraet/facebook\\_belgium](https://github.com/SPraet/facebook_belgium)

and demonstrate the potential added value of combining the predicted Facebook Likes models with traditional survey data.

More specifically, we compare the predictive performance of Facebook Like data to the survey data in four different set-ups: models based on all Facebook Likes (M1), models based solely on non-political Facebook Likes (M2), models based on the survey data (M3), and models based on a combination of all Facebook data and survey data (M4) (see Table 5.2). Prediction models are evaluated based on out-of-sample predictive accuracy. We report the average AUC over ten folds.<sup>10</sup>

Political leaning is divided into three classes (left, center, and right). For each of the three classes, we transfer the variable into a binary classification problem (one-vs-all) and train three binary Logistic Regression (LR) models.<sup>1112</sup> Next, we calculate the weighted average AUC for the three classes. Similarly, we built seven binary classifiers for the seven political parties and calculate weighted average AUC. Though multinomial LR (one-vs-one) could also be used in this case (and is in fact more common in electoral research), we prefer binary classification to be able to create general distinct profiles for each party electorate in comparison to all other citizens, rather than compare the voter profiles to only one single reference group. However, both methods (binary and multinomial LR) show similar predictive performance in our analysis.

Additionally, we want to know how predictive each category of pages is for political leaning. In other words, how accurate can political leaning be predicted when using only the Facebook Likes of the concerning category. This tells us which aspects of our social lives are most related to politics. To report the predictive performance of each category independent of the amount of pages per category, we randomly sample (with replacement) 100 pages per category and use only those pages as features.<sup>13</sup> This procedure (of random sampling) is repeated 10 times, and the average AUC is reported.

Finally, we compare the *insights* in ideological position and party preference based on Facebook Likes versus survey data. In Appendix 9.3.2 we compare six methods to gain insights from sparse data and we suggest to rank the Facebook pages based on the coefficients of a regularized logistic regression. Analyzing the coefficients of a logistic regression and their “p-values” is common practice in traditional explanatory and predictive modeling with dense data (Francis and Payne, 1977). However,

<sup>10</sup> This means that we train the model on 90% of the data and evaluate on an unseen 10% of the data, this procedure is repeated ten times with different parts of the data. For an explanation of AUC see Chapter 2.

<sup>11</sup> We used the scikit-learn implementation for logistic regression (Pedregosa et al., 2011b)

<sup>12</sup> The classifiers were trained using 5-fold cross validation to optimize the regularization penalty ( $L1$  or  $L2$ ) and the optimal regularization value  $C$  in  $[0.001, 0.01, 0.1, 1, 10]$ . The data is split into five folds, for the different parameters a logistic regression is trained on four folds and the parameters that resulted in the highest predictive performance on the test fold are selected.

<sup>13</sup> Because we do not have all categories of the less frequent pages available, we will only include pages with 30 likes or more.

Table 5.2: Overview of the modeling set-ups with a description of the different datasets and the number of variables. In the description we refer to the survey questions in Appendix 9.3.1.2.

Data	Description	# Variables
Facebook Likes (M1)	Public Facebook Page Likes	595,994
Facebook Likes* (M2)	Non-political public Facebook Page Likes, liked more than 30 times	10,226
Survey (M3)	Use of media sources (Q1), interest in news topics (Q2), interest in politics (Q5) Gender (Q6), age (Q7), education level (Q8)	27
Target variables	Political leaning (Q3) and party preference (Q4)	10

when using a regularization term, assumptions about the asymptotic distribution of parameters do not apply, and therefore, different methods for significance testing are needed and suggested in literature (Lockhart et al., 2014; Tibshirani et al., 2015). We will follow the bootstrap procedure as described in Tibshirani et al. (2015). From the original dataset we take a random sample with replacement and built a model from this dataset to estimate the coefficients. This step is repeated 1,000 times to obtain 1,000 values for each coefficient. For each coefficient we estimate the probability density function using a Gaussian kernel<sup>14</sup> to calculate the probability (p-value) that the parameter is less than or equal to zero. Facebook pages will be ranked based on the mean coefficients over 1,000 bootstraps while the p-value indicates significance on a  $\alpha = 0.05$  level.

Alternatively, instead of analyzing individual Facebook pages, we can apply dimensionality reduction to group related pages together. We apply Non-negative Matrix Factorization (NMF)<sup>15</sup> to transform the data from a sparse representation to a dense one. The matrix does no longer exist of one-hot encoded vectors but each user receives a continuous score (between 0 and 1) for every dimension. Similar to the procedure described above, we built a logistic regression model on the reduced dimensions and investigate the coefficients of this model. The number of topics ( $k \in \{50, 100, 150, \dots, 400\}$ ) is optimized based on the performance of the subsequent classification task. This way, the number of topics is set to  $k = 200$ .

<sup>14</sup> We used `gaussian_kde` from `scipy.stats` (Jones et al., 2001) with the default Scott's Rule (Scott, 2015) for bandwidth selection.

<sup>15</sup> See Chapter 2

## 5.5 Results

We start this section with an evaluation of the models to predict political leaning and party preference. Secondly, we look at the political insights that can be gained from these models.

### 5.5.1 Predictive performance

Table 5.3 compares the predictive performance for models based on Facebook Likes and survey data for all users that participated in our study. We use a corrected paired differences t-test (based on 10-fold cross-validation results) to compare the AUC of the different models (Nadeau and Bengio, 2000) and apply the Bonferonni correction for multiple pairwise comparison (Vázquez et al., 2001; Pizarro et al., 2002). For ideological leaning, all pairwise differences between the models are significant at the  $\alpha = 0.05$  level<sup>16</sup> (See Table 9.15 in the Appendix). The Facebook Likes (M1) exhibit a higher predictive power than the survey data (M3). Still, the survey questions do seem to capture some information that Facebook Likes do not, since the combined model (M4) achieves the highest predictive power. One might argue that the high predictive performance of Facebook Likes is due to the presence of political Facebook pages. Naturally, if someone likes the Facebook page of a certain party or politician this is very indicative of their political preference. Therefore, to assess the predictive value present in non-political Facebook pages, we built the same prediction models on the Facebook data without political pages (M2). As expected, the predictive performance decreases when excluding explicit political content. However, the AUC is still higher than the survey data (M3).

The high predictive performance of non-political Facebook Likes is an indication that non-political interests are interrelated with political preferences. The fact that this performance is higher than for the available survey data indicates that indeed additional information is captured with Facebook data that was not covered in the survey questions. To further explore which lifestyle categories are connected to political preference, we calculate the predictive performance per Facebook category. Figure 5.1 shows which categories of Facebook pages are most predictive for ideological preference. A prediction model built on only political Facebook pages achieves, not surprisingly, the highest AUC, followed closely by **Civil Society** and **News & Media**, which could be considered as semi-political categories. However, also the non-political categories of **Arts & Culture** and **Communities** are very telling for ideological preference. Conversely, the categories **Movies**, **Sports & Health** and **Travel** are the least predictive for ideological preference. This indicates that some aspects of our social lives are more tied to our ideological views than others.

Finally, we note the non-political likes perform better at discriminating between left versus non-left people and between right versus non-right than between center

<sup>16</sup> with the Bonferonni correction this is reduced to  $\alpha/6$

Table 5.3: AUC and the standard deviation (std) for models built on (non-political) Facebook pages, survey data or a combination thereof.

		Facebook (M1)		Facebook* (M2)		Survey (M3)		Combined (M4)	
		AUC	Std	AUC	Std	AUC	Std	AUC	Std
Leaning	Left	83%	2%	75%	3%	71%	3%	84%	2%
	Center	64%	2%	62%	3%	62%	2%	66%	2%
	Right	81%	3%	75%	3%	70%	4%	81%	3%
	Weighted average	73%	2%	68%	2%	65%	2%	74%	2%
Party preference	PVDA	79%	3%	73%	3%	66%	3%	79%	2%
	Sp.a	70%	4%	70%	2%	61%	2%	69%	3%
	Groen	79%	1%	71%	1%	69%	1%	79%	1%
	CD&V	73%	3%	68%	3%	58%	3%	72%	2%
	Open VLD	77%	4%	73%	3%	67%	3%	77%	4%
	N-VA	84%	3%	80%	2%	70%	2%	84%	3%
	Vlaams Belang	86%	5%	76%	7%	76%	7%	85%	7%
	Weighted average	78%	1%	73%	1%	67%	1%	78%	1%




Figure 5.1: Average AUC (+/- 1 std) for the Facebook categories with target "left".

versus non-center (see Table 9.15 for the t-test results). This might indicate that a less clear pattern is present in the characteristics and behavior of people with a center political leaning and that they exhibit a less pronounced profile. Similarly, voters for some parties can be classified more accurately than others (e.g. the party N-VA, see Table 9.16).

### 5.5.2 Insights

From the survey data (see Table 9.18), we learn that traditional variables, such as gender, age, education, and interest in certain topics help to predict people's vote. For instance, left voters are often highly educated women, between the age of 25 and 55, and interested in news about culture, arts, and international politics. Conversely,

right voters are more often male than female and generally demonstrate a strong interest for financial and economic news. What additional insights do Facebook Likes provide? To answer this question, we examine the most related pages when excluding political pages (see Table 9.17 and Figure 5.2 for a summary).

With Facebook Likes, analyzing cultural taste and differences in lifestyle between left and right voters is possible at a very detailed level. The pages most related to a left political leaning are (alternative) media outlets that are considered to be leaning more toward the left, and nonprofit organizations for climate and human rights, which can be considered as outspoken left-wing topics. For example, all else equal, the odds of demonstrating a left political ideology are almost 40% higher for someone who liked the website of the left-leaning newspaper *De Morgen* than for who did not, and they are almost 20% higher for someone who liked the nonprofit organization *Amnesty International*. Right pages are dominated by Flemish nationalistic content and memes. The odds of demonstrating a right political leaning are almost 20% higher for someone who liked the alternative right-wing news website *SCEPTR* or the popular mainstream paper *HLN.be* than for someone who did not. Similarly to the left pages, most of these pages carry a subtle, or sometimes an outspoken, reference to political ideology. In contrast, center voters like less explicit or implicit political pages on Facebook.

Similarly, to predict political parties, politically loaded pages are ranked high, such as newspapers or organizations with a certain ideology. Bluntly summarized, voters for the workers' - and social democratic party like solidarity content such as refugee or third-world organizations, green voters mainly like environmental pages (e.g. Greenpeace) Christian democratic voters are interested in religious organizations and the royal family, liberal voters in financial news and businesses, and finally the Flemish nationalists and extreme rightists tend to like Flemish nationalistic and identity content on Facebook. At first sight, most of these likes seem to be in line with the broader ideological or issue profile of the party family.

When delving deeper into aspects of lifestyle and cultural preferences our methodology allows us to focus on the most related pages per specific Facebook category. For example, when considering only the pages in the category **Movies**, left voters more often like adventure, romance and drama movies, such as *The Hunger Games* (Odds Ratio (OR) = 1.06) or *500 Days of Summer* (OR = 1.06), and visit arthouse cinemas in larger cities. In contrast, action movies such as *Scarface* (OR = 1.04) and *Fast & Furious* (OR = 1.04) are mainly liked by voters on the right. In the category **Music**, the genres of alternative rock, blues, and experimental music are most related to a left political leaning (e.g. *Bob Dylan* (OR = 1.08) or *Tom Waits* (OR = 1.07)) whereas popular (hard) rock music (e.g. *AC/DC* (OR = 1.05)) is often liked by right voters, next to the genres of techno and electronic dance music (e.g. *Justice* (OR=1.04)). The same analysis can be done for other categories such as books, food, brands, sports, etc. Clearly, the lower odds ratios of a single movie or rock band compared

		Left	Center	Right
Survey	Socio-demo (Q6, Q7, Q8)	Female, higher educated, 25-55	Male, -25 or 55+	Male
	News subjects (Q2)	International, arts & culture	Finance & economy, international, justice & safety	Finance & economy, justice & safety
Facebook	Likes	Left media, nonprofits, art & culture	Financial news, entertainment, food	Right media, financial news, Flemish nationalistic content
	Companies	Environment, healthy food, local	Retail, travel	Financial, cars, beer
	Movies	Adventure, romance, drama	Family movies	Action and crime movies
	Music	Alternative rock, blues, experimental	Belgian and pop music	Popular rock, techno, electronic

Figure 5.2: Some examples to illustrate the insights we can gain into voter profiles based on Facebook data and based on survey data.

to those of a civil-society organization or alternative news outlet indicate that their connection with ideological preference is present, but much more modest.

An alternative way to gain insight in voter profiles is through dimensionality reduction. We'll illustrate this with an example for the left and right target class. A logistic regression based on the NMF components to predict a left leaning achieves an AUC of 78%. Table 9.19 in the Appendix shows the most related latent dimensions to a left ideological leaning. The NMF components seem to group the Facebook pages in an intuitive way which facilitates the interpretation. One component is largely related to the green party (environmental concerns) and one to the workers' party (solidarity) but some overlap between both components exists. Furthermore, cultural attractions in the city of Ghent and Antwerp are grouped together. The fifth component contains mainstream media and some alternative media that are considered to be leaning towards the left (i.e. *De Wereld Morgen*, *Apache* and *MO\**).

The most related components to a right political leaning can be found in Table 9.20 in the Appendix. The first and second component group politicians and other pages related to the Flemish nationalists party and the liberal party respectively. The third component groups local branches of the Flemish nationalists party. The fourth component consists of soccer players and finally, the fifth component contains (electronic dance) music. This technique has the advantage that several related Facebook pages are automatically grouped together in logical components. These groups of pages are more telling than single Facebook pages. For example, the music component for right political leaning provides a more elaborate view on music taste compared to a single Facebook page of a band or artist.



## 5.6 Conclusion and future research

The starting point of this study is that lifestyle and politics are closely related. Next to classical variables such as socio-demographics and issue preferences, our personal values and lifestyle choices correlate with our political preferences. The fact that (non-political) Facebook Likes achieve high predictive accuracy (in addition to survey data) shows that they are indeed capturing additional information, which we argue are related to values and lifestyle.

Consequently, we looked into which aspects of our social life are most predictive for our political preference by analyzing the predictive performance of Facebook categories (e.g. movies, music, food, etc.) separately. Non-surprisingly, politicians, media and civil society are most predictive for political leaning, but also arts, culture, entertainment and books help to predict where people stand politically, whereas, in particular sports and travel are less predictive. This raises the question why some aspects of our social lives are more connected to political preference than others. For instance, our study seems to suggest that Belgian citizens are more likely to meet people with different political convictions when cheering for their favorite sports team than when going to a music festival. However, the extent to which these results can be generalized to other contexts or countries is uncertain. For instance, the study of Shi et al. (2017b) found that for Twitter users in the U.S. musical preference was less connected to political preference than following sports teams. In the context of an authoritarian state, Urman (2017) found that several non-political interests (e.g. rappers, history, travel) were indicative of pro-opposition partisanship on the Russian social network site Vkontakte. To reach insights that travel across countries and time periods, future research needs to be comparative or at least more similar in terms of data and methods.

In contrast with most previous research on lifestyle politics in the (polarized) two-party system in the U.S., we examined lifestyle politics in a multi-party system, with much more subtle ideological differences between parties. Our study indicates that Facebook Likes are less predictive for center voters and for traditional political parties. A less clear pattern is present in the characteristics and behavior of those voters compared to voters with a more outspoken ideological position. We find, for instance, that our social media data are much better in predicting who votes for an extreme-right party compared to the social democrats. Is this simply due to the more radical or straightforward ideological position of these parties, or do certain politicians link their ideology or party platform more to lifestyle choices? How to explain these differences and why certain ideological and political opinions are more connected to lifestyle than others, are interesting follow-up questions for social scientists, and electoral scholars in multi-party systems in particular.

Our analyses showed that with Facebook data, different interest categories can easily be analyzed and compared to improve our understanding of public opinion and voter behavior. Capturing this amount of detailed information using traditional

surveys is very difficult, as it would imply long question batteries on different categories related to lifestyle and consumer choices (but see DelaPosta et al., 2015). In addition, it would request scholars to define a priori the most important political lifestyle indicators. An important aspect of future research is then how to summarize the insights gained from Facebook Likes in such a way that it can support theory building. Analyzing the most predictive Facebook pages per Facebook category is one way, and has the advantage that we can investigate the categories of interest in a structured way. Another possibility is dimensionality reduction to learn latent dimensions from the data itself. For instance, we found that music of the same genre was grouped together. These dimensions can help us to identify overarching cultural taste or lifestyle patterns which can thereafter be transformed into a series of survey questions and used for explanatory analysis.

At the same time, we do not argue that social media data can replace survey research. Probably both advanced survey research and social media data are needed to understand how lifestyle and political preferences influence each other. For instance, DellaPosta and colleagues (2015) explain the puzzling association between lifestyle preferences and political affiliation by the self-reinforcing effect of homophily and social influence. In short, they argue that people make lifestyle choices and opt for places where they meet likeminded people (homophily), in those (online) places people's attitudes are affected by exposure to each other (social influence). However, to better understand how this self-reinforcing effect takes place, further research is needed. By using a survey panel design (Dvir-Gvirsman, 2017) or by simultaneously analyzing friend networks on Facebook (Bond et al., 2012), the relationships between homophily and social influence could be further explored. Interviews or surveys could seek for explanations for this association and its direction by asking questions such as "Do you like to go to events where people have the same opinions as you?" or "Do you talk about politics or societal issues at music concerts?".

We argue that combing social media and survey data is not only useful to understand where "the cultural fault lines" are most persistent, but it can simultaneously provide insights in which (type of) organizations, brands or events "unite" people with different political views. For example, our data indicate that people from all sides of the ideological spectrum "like" a political comedy show from the public broadcaster (*De Ideale Wereld*). From the perspective of the recently growing polarization of the audience for late night comedy in the U.S. (Young, 2019), this is not a trivial finding. Therefore, this type of knowledge may be increasingly relevant and useful, as ever more countries are facing increasing levels of (affective) polarization (Iyengar et al., 2012; Lelkes, 2016). Furthermore, the growing affective polarization in society might make people more eager to express their political identity online by liking pages or following actors that are seen in line with their political views and cultural tastes. This would imply that similar factors drive both online and offline lifestyle choices. Further research could tackle this, among others by studying the linkage between social media data and political preferences over time.

In conclusion, we argue that fine-grained behavioral data are useful to discover unknown patterns or better understand existing relationships and are therefore equally valuable for political and social science research. However, the mining of behavioral and online data has raised new and unexpected ethical and regulatory questions (Greene et al., 2019). We therefore underline the importance of an ethical discussion of the research design, which we included in Chapter 2. Unfortunately, Facebook restricted data collection through the APIs of Facebook, Instagram, and other platforms it owns. While this intervention certainly is positive for the privacy protection of their users, it is also locking out third parties and diminishing transparency of the platform. Nonetheless, the potential of other social media data (e.g. Twitter) or other types of behavioral data (e.g. location data, payment data or browsing data) could be further explored as they complement the insights of survey research into different aspects of citizens' social and political life.



## Polarization in lifestyle domains

---

Increasing levels of political animosity in the United States invite speculation about whether polarization extends to aspects of daily life. However, empirical study about the relationship between political ideologies and lifestyle choices is limited by a lack of comprehensive data. In this research, we combine survey and Facebook Like data from more than 1,200 respondents to test whether polarization permeates society or if it is more limited to strictly political domains among politically active individuals. Our results indicate that polarization is present in page categories that are somewhat related to politics — such as opinion leaders, political news sources, and topics related to identity and religion — but, perhaps surprisingly, it does not appear to have strong influence in other domains, including sports, food, and music. On the individual level, we find that people who are higher in political news interest, have stronger ideological predispositions, and are over the age of 65 are especially likely to endorse ideologically homogeneous pages across categories. Our evidence, drawn from rare digital trace data covering more than 5,000 pages, adds nuance to the narrative of widespread polarization across lifestyle sectors, and it suggests domains in which cross-cutting preferences are still observed in American life.

## 6.1 Introduction

During the last few U.S. presidential election campaigns, figures across diverse sectors of society used their platforms on social media to persuade or mobilize their fans. For inattentive or first-time voters, this may have offered a rare encounter with political advocacy as well as the authenticity that can be gained when mostly apolitical actors or entities take a stand (Zilinsky et al., 2020). Fans of L.A. Lakers basketball player LeBron James (with over 23 million followers on Facebook) may be familiar with this dynamic — as well as its pitfalls. As James took a more prominent role supporting voter registration efforts and criticizing President Trump in the midst of the pandemic and a summer of roiling protests over racial justice, he found himself the target of attacks by the president himself. In a period of intense affective polarization at the mass level, politics threatened to ensnare one of the most popular celebrities in America.

Athletes and sports leagues have again found themselves at the crosshairs of political controversy in the United States, as symbolic demonstrations of racial solidarity in the form of kneeling protests have become increasingly commonplace, sometimes revealing divides between players and their fans. These episodes are an especially vivid illustration of how seemingly apolitical domains — including sports but also food, artistic and cultural preferences, and consumer decisions — can become caught in the partisan currents of the larger society. More specifically for the focus of this study, they also demonstrate the importance of social media as an arena where lifestyle preferences in all their dimensions can intersect with politics. Given growing concern that political polarization is permeating society, this study asks *to what extent* this phenomenon is reflected in other realms. Widespread polarization across lifestyle domains would have serious implications because cross-cutting pressures in formally apolitical spheres may be critical for maintaining social harmony in otherwise highly polarized political systems (Mutz, 2006; Pettigrew, 1998).

In the previous chapter, we showed how Facebook Page Likes can be used to study lifestyle preferences and their link with political affiliation. In this study, we take advantage of the comprehensive revealed preference information from Page Likes to understand whether preference sorting across various lifestyle categories follows the pattern established in partisan politics. By combining survey and Facebook Likes data from more than 1,200 respondents in the United States — Facebook is used by nearly 70% of Americans<sup>1</sup> —, we directly test whether pages belonging to more “political” categories will be liked by more polarized audiences and whether individual-level characteristics are associated with liking pages in more polarized categories. We find, in contrast to some existing work, a clear divide between more polarized page categories that are commonly understood to be related to politics — such as opinion leaders, political news sources, and topics related to identity and religion — and other categories, including sports, with relatively low levels of

<sup>1</sup> Source: <https://www.pewresearch.org/fact-tank/2019/05/16/facts-about-americans-and-facebook/>

ideological homogeneity within pages. On the individual level, we find that people with higher political news interest, stronger ideological predispositions, and who are over the age of 65 are especially likely to endorse such ideologically homogeneous pages across categories.

## 6.2 Polarization, lifestyle, and social media

There is a longstanding debate about the nature and extent of polarization in American society, defined as party differences in issue positions or attitudes (Evans et al., 2001; Fiorina et al., 2008; Hetherington, 2009; Baldassarri and Gelman, 2008; Abramowitz and Saunders, 2008). To synthesize this large body of research, it can be argued that today's high level of observed partisan polarization in ideological affiliations and issue positions is largely a reflection of increased sorting among the parties over time (e.g., conservatives into the Republican Party and liberals into the Democratic Party), and that evidence of increasing extremity over time is mixed, especially among the mass public (e.g., Lelkes, 2016). More recently, scholarly attention has shifted to the affective dimension of polarization, rooted in an understanding of partisanship as a social identity (Mason, 2018b; Iyengar et al., 2019; Finkel et al., 2020). This aspect of polarization, distinct from specific attitudes, is particularly important for understanding how personal and emotional attachments formed in the political arena could potentially carry over to other domains. Such a process is suggested by the "oil spill" model of polarization, in which clusters of initially disparate issues — including cultural and moral issues — become connected in a belief system (DellaPosta, 2020). The potential for polarization to spread beyond strictly political settings is also suggested in a conception of partisan attachment as reflecting a shared understanding of constituent social groups (Green et al., 2004). Related conceptions have likewise been proposed for understanding ideologies as a basis for group identification (Conover and Feldman, 1981; Mason, 2018a).

The literature on assortativity documents how political identities have begun to structure social behavior in everything from online dating behavior (Huber and Malhotra, 2017) to marriage (Alford et al., 2011; Iyengar et al., 2012), while debate continues on the extent to which people choose to locate geographically in areas with like-minded partisans (e.g., Bishop, 2009; Mummolo and Nall, 2017). Preferences likely extend to entertainment choices such as television viewing habits as well (Toff, Nd). It is important to note that while average partisan differences have been clearly documented in these areas, gaps can be exaggerated relative to overall levels. Such a dynamic has been observed in persistent exaggerations of partisan differences in demographic composition and other common partisan stereotypes, for example (e.g., Ahler and Sood, 2018).

Research similarly suggests a nuanced understanding of how social media reflects these patterns. A large study of retweet networks found much more ideological homogeneity between users tweeting about political topics than about nonpolitical

topics, such as the 2013 Boston Marathon bombing in its initial aftermath and the 2014 Super Bowl (Barberá et al., 2015). But the marathon bombings themselves became politicized over time, and polarization in these retweet networks increased as a result. How topics can come to be seen as “political” or not is itself a challenging question, as Settle (2018) argues in the context of Facebook’s News Feed. In Settle’s theory, politically inattentive Facebook users come to make inferences about their more political friends by observing their posts and endorsements (including via likes). Through this process, associations come to form between political identities and lifestyle preferences. This is vividly illustrated in the book by the example of Chick Fil-A, which became a flashpoint in America’s culture wars over the issue of same-sex marriage, which the chain’s owner publicly opposed. To some, the choice of fast-food chain for a quick meal may not reflect political preferences, but such decisions can nonetheless take on a larger symbolic meaning to outside observers.

Settle’s argument raises the question of whether social media, and Facebook specifically, is accelerating the process of politicizing lifestyle choices and preferences so that they more closely map onto the partisan political divide. Since our data provide a snapshot in time, this study cannot specifically answer this question, though it sheds light on the baseline levels of polarization across different areas of society. However, evidence is accumulating for the specific mechanisms likely at play, namely inferences due to apolitical cues (e.g., Lee, 2020).

### 6.3 Hypothesis and research questions

A large amount of political behavior research on social media, including Twitter (Barberá, 2015; Boutyline and Willer, 2017; Eady et al., 2019) and Facebook (Bakshy et al., 2015; Bond and Messing, 2015), focuses on retweeting, following or liking of political figures, news media or other political content. Despite Facebook’s importance for our understanding of social media and political communication, average users of the social platform likely do not primarily use it to follow news or engage in politics.

We analyze Facebook Likes to gain insights into which aspects of people’s social lives appear to reflect strong ideological divides. We construct measures of the ideological *homogeneity* of users who like individual pages, allowing us in turn to draw conclusions about the overall polarization of liking patterns across “political” and “non-political” categories. At the page level, we first ask: *How ideologically homogeneous are political and lifestyle categories on Facebook? (RQ1)* We expect political pages to exhibit relatively high levels of ideological homogeneity since online endorsement of political figures is a reflection of political ideology and opinion (Bond and Messing, 2015). Prior research has found that motivations to Like a political page include self-reflection, showing political interest, and seeking engagement with others (Macafee, 2013). Although there are various reasons to refrain from liking the Facebook pages of political candidates (such as social anxiety and audience diversity;



see Marder 2018), it is relatively uncommon to like the Facebook page of an opposing candidate or party.

We may similarly expect divides in the audiences of news sources' Facebook pages. The most recent evidence using behavioral data on individuals' online media consumption suggests nontrivial overlap in Democrats' and Republicans' news diets, which consist to a large extent of relatively centrist mainstream outlets and large portals (Guess, 2020). Data on people's follow networks on Twitter similarly suggest a meaningful amount of exposure to cross-cutting news (Eady et al., 2019). Twitter allows users to follow accounts without this necessarily being visible to the public (e.g., by creating a "private" list). However, the more observable nature of people's page endorsements on Facebook suggests a potentially greater role for political identities to shape self-presentation, which could affect conscious decisions to like sources perceived as politically congenial. In this way, we expect patterns of Facebook news likes to more closely mirror survey evidence (e.g., Newman et al., 2019a; Jurkowitz et al., 2020), which scholars have argued may reflect partisan biases in addition to (or in place of) accurate reporting of news consumption habits (Prior, 2013).

Finally, even non-political lifestyle categories (arts and culture, food, sports, etc.) have been found in studies based on survey data to map onto ideological divides (DellaPosta et al., 2015; Hetherington and Weiler, 2018), though we would expect ideological diversity to be higher for lifestyle domains than for more explicitly political categories. To the best of our knowledge, the only social media study that has explicitly focused on the partisan divide in non-political domains in the United States was conducted by Shi et al. (2017b), who analyze Twitter co-following networks.<sup>2</sup>

Next, we turn to the individual level and ask whether traditional predictors of strong political engagement (e.g., Saunders and Abramowitz, 2004) also predict liking relatively polarized *non-political* pages. This would indicate that the predictors of polarized endorsements of traditionally political content carry over into lifestyle and related domains. Therefore, we propose the following hypothesis: *People with greater political interest and stronger ideological affiliations are more likely to like politically homogeneous Facebook pages in non-political categories (H1)*. Thus for both political and non-political pages we would expect "very liberal" and "very conservative" ideological affiliations and high levels of political news interest to be significantly related to high page homogeneity scores, all else equal.

To further understand how individual-level characteristics relate to preferences for lifestyles and habits shared among relatively homogeneous groups, we build upon research from political psychology to explore which other characteristics are associ-

<sup>2</sup> Facebook conducted an informal, proprietary analysis in 2014 focusing on the music, TV, and other cultural preferences of users who liked official political pages during the midterm election campaign. See <https://www.facebook.com/notes/facebook-data-science/politics-and-culture-on-facebook-in-the-2014-midterm-elections/10152598396348859/>.

ated with higher levels of political homophily,<sup>3</sup> and ask a second research question: *Is ideology associated with a greater likelihood of “liking” ideologically homogeneous pages?* (RQ2) We will focus on differences between political pages, news and media pages, and lifestyle pages. Jost et al. (2009) argue that conservatives and liberals differ in their need for certainty, making intolerance of ambiguity more typical of the political right as compared to liberals’ greater openness to new experiences and cognitive complexity. On the other side of the asymmetry debate, Greenberg and Jonas (2003) show that individuals on *either* ideological extreme possess greater preference for certainty than more moderate ones. At the same time, there is some evidence that liberals prefer more homogeneous content: Bakshy et al. (2015) find that liberals tend to be connected to fewer ideologically dissimilar friends on Facebook than conservatives, while Eady et al. (2019) find that liberals are less likely to follow media and political accounts classified as right-leaning than vice versa.

Based on this inconclusive body of research, we tentatively explore ideological self-placement as a possible characteristic associated with liking ideologically homogeneous pages. In a similar spirit, we are interested in potential age-related differences as well: *Is age associated with a greater likelihood of “liking” ideologically homogeneous pages?* (RQ3). Recent research exploring the determinants of “fake news” sharing on Facebook suggests a strong age effect (Guess et al., 2019b). Though that study did not directly test the extent to which the likelihood of sharing online misinformation is related to its prevalence in users’ News Feeds (e.g., see Guess et al., 2020), such a link would suggest a role for liking pages associated with untrustworthy, highly polarized posts.

#### 6.4 Data collection

A panel survey on (social) media use was conducted during the 2016 U.S. presidential election ( $N = 3,500$ ) over three waves. The respondents were asked to complete a survey and to indicate their ideology on a 5-point scale (see results in Table 6.1). In Wave 1 (April 9–May 1, 2016), we also asked respondents to indicate social networking sites for which they had accounts (options: Twitter, Facebook, Instagram, LinkedIn, Snapchat, and other).

After the election we asked respondents if they would be willing to supply information about their own past Facebook activity. This was done via a separate survey question that sent respondents to a web application facilitating an authenticated link to the Facebook API. 1,331 respondents consented to let us retrieve their Facebook information.<sup>4</sup> Specifically, we requested their public profile information, Timeline posts (including text and links if available), Page Likes, and what Facebook saves

<sup>3</sup> Political homophily refers to the tendency to associate with others who are similar in political ideology (for a review, see Boutyline and Willer, 2017).

<sup>4</sup> Respondents who consented to provide their Facebook data were compensated with an additional \$5 in YouGov “points” above what they received for taking the survey

Table 6.1: Sample details

Ideology	Respondents
Very liberal	189 (16%)
Liberal	218 (18%)
Moderate	360 (30%)
Conservative	175 (14%)
Very conservative	92 (8%)
Not sure	51 (4%)
No answer	126 (10%)

as religious and political views. If a respondent chose to log into Facebook after the survey prompt, they were asked what specific pieces of information they were willing to share. They could approve sharing all of the given types of information, selectively approve only some of these types of information, or approve none of them.<sup>5</sup> Of the 1,331 respondents (comprising 45% of the 2,711 respondents who reported having a Facebook account) who agreed to share Facebook profile data, 1,230 could be successfully linked back to the survey for our study. Based on observable characteristics, the subgroup for which we have Facebook data is a fairly representative cross section of the overall sample (see Table 9.21 in Appendix 9.4.1). However, those who shared data were slightly more liberal on average, more likely to participate in elections, and more politically engaged.

Our dataset consists of 387,671 unique Facebook pages. The majority of these pages are liked by fewer than 5 respondents in our dataset (see Figure 6.1). To ensure that our results are not being impacted by small numbers of people liking particular pages, we restrict our analysis to pages that are liked by at least 30 respondents each.<sup>6</sup> Two independent coders were trained to categorize all Facebook pages into predefined categories (e.g. politics, news, sports, or food).<sup>7</sup> Since the categories are not mutually exclusive (for example, LeBron James is included in **Sports** as well as **Public Figures**), the coders could assign a maximum of 3 categories per page. The coders agreed on at least one category for 70% of the Facebook pages. The categories and their description can be found in Table 6.2.

<sup>5</sup> No data on News Feed content or exposure was shared with researchers. Data access was temporary and lasted only 2 months after permission was granted. All respondents who agreed to share information consented to a privacy policy that specified, in part, “This application will not access the profile information of any friends, groups, or other information associated with your profile page.”

<sup>6</sup> This way, the number of pages included in our analysis is reduced to 5,155, still accounting for almost 25% of total Page Likes. In Appendix 9.4.4.1, we show that the individual-level Like distributions of all pages and pages with a minimum of 30 likes are highly similar. We acknowledge that the results and conclusions in this study are based on relatively popular Facebook pages and that we cannot analyze polarization on smaller pages.

<sup>7</sup> Details on the coding task can be found in Appendix 9.4.2




Figure 6.1: Page Likes distribution

Table 6.2: Description of the Facebook categories and the number of pages per category

Category	Description	# pages
Shopping & retail	Apparel, accessories, clothing, fashion, consumer electronics, home decoration, stores, shopping mall, wholesale, etc.	1858
Public Figures	Public figures	876
Food & Beverage	Food, cooking, restaurants, drinks, spirits, breweries etc.	837
Entertainment	Entertainment, games, humor, amusement, comedy etc.	532
Music	Music, bands, producers, record labels, albums, awards, concerts, music festivals etc.	499
Tv Shows	TV shows, episodes, channels, TV awards	463
Politics	Politicians, political parties, political content, political communities and government organizations	353
Movies	Movies, actors, directors, movie characters, cinema and awards	341
Services	Marketing, advertising, legal, finance, consulting, etc.	339
Beauty & Health	Cosmetics, healthcare, medical	337
Civil Society	Nonprofit organizations and labor unions (formal organizations)	211
Political news	News and media about politics	188
Interests	Interests, communities (informal) and hobbies	166
Arts & Culture	Arts, culture, photography, museums, artists, musicals, theater, literature, libraries, writers, etc.	110
Hard news	Factual reportage of events which are socially or politically significant and of a serious nature	103
Sports	Sports, teams, athletes, leagues, games, gym	99
Cars and transportation	Car brands, automotive, airlines, boats, etc.	74
Identity & Religion	Pages referring to home country, region, ethnic or cultural groups, religious pages, religious organizations	72
Travel	Travel, tour agencies and tourism	70
Individual opinion leaders	Individual influencers, bloggers, commentators, etc.	51
Research & Education	Schools, universities, student organizations, educational programs, (non-)scientific research	43
TOTAL		5155

At the highest level, we distinguish between three groups: pages in the category of (1) Politics, (2) News & Media, and (3) all other “Lifestyle” categories.<sup>8</sup> For these three groups the percentage agreement between two coders is 90%. We further break the second group down to (2a) Political news and (2b) Hard news. This was done manually by the authors. “Hard news” refers to non-partisan news sources. “Political news” refers to partisan news sources such as HuffPost. News outlets that combine “hard news” with clearly partisan opinion desks (e.g., Fox News) are included in both categories. News categories related to lifestyle, celebrities, sports, and science are included in group (3) Lifestyle.

## 6.5 Methods

We start with network analysis and community detection to study the extent of political polarization on Facebook. We first create a bipartite network in which Facebook users are the bottom nodes and Facebook pages are the top nodes. An edge exists between a user and a page when the user has liked the page. Next, we project the bipartite network to a homogeneous unigraph of the bottom nodes, where Facebook users are linked if they have liked a common Facebook page (see Figure 6.2).




Figure 6.2: Bipartite network with Like relationships between users (bottom nodes) and Facebook pages (top nodes), and the top- and bottom-node projections.

The weight of the edge between each pair of nodes captures the number of shared Facebook pages. The Louvain algorithm (Blondel et al., 2008) is applied on this bottom node projection to detect communities of highly connected users. The algorithm is based on modularity optimization. Modularity (see Equation 9.9 in Appendix 9.4.3) is the relative density of edges inside communities with respect to edges outside of communities and measures the extent to which a network is divided into different clusters or communities. Networks with high modularity have dense connections between the nodes within clusters but sparse connections between nodes

<sup>8</sup> Civil Society, Public Figures, Individual opinion leaders, Research & Education Arts & Culture, Tv Shows, Entertainment, Movies, Interests, Music, Sports, Beauty & Health, Food & Beverage, Shopping & retail, Travel, Cars and transportation.

in different clusters. In practice, a modularity value above 0.3 is a good indicator of significant community structure in the network (Clauset et al., 2004). The existence of ideologically different communities would indicate polarization in the network, since ideologically similar individuals are more strongly connected to each other than to ideologically dissimilar ones. Separately, we perform a top node projection of the bipartite network, where the weights of the edges between Facebook pages reflects the audience overlap, and we repeat the analysis.

Next, we analyze the ideological homogeneity<sup>9</sup> of different lifestyle categories and users on Facebook in more detail. We measure the ideology and homogeneity of individual Facebook pages based on the liking behavior of our respondents and average this over categories and individuals. In the following discussion, consider Facebook page  $Z$ , self-reported individual ideology score  $k \in \{0, \dots, 4\}$  (where  $0 = \textit{very liberal}$  and  $4 = \textit{very conservative}$ ), and ideology class  $c \in \{0, 1, 2\}$  that groups these ideology scores  $k$  into three groups (where  $0 = \textit{liberal}$ ,  $1 = \textit{moderate}$ , and  $2 = \textit{conservative}$ ).

**MEASURING PAGE IDEOLOGY** Using like behavior and self-reported ideology of the respondents in our sample, we map the ideologies of Facebook pages. Several behavioral approaches can be found in the literature. For example, Bakshy et al. (2015) estimate the ideology of news media by calculating the difference in the proportion of self-reported liberals and conservatives who share links to such media on Facebook, while Messing et al. (2017) calculate media ideology by averaging the NOMINATE scores of members of Congress who share news media URLs. Similar to these approaches, we average the self-reported ideology scores ( $k$ ) of respondents who liked Facebook page  $Z$  to calculate the page ideology score of Facebook page  $Z$ , which ranges from 0 to 1. To adjust for uneven partisan distribution, we add a correction factor of 0.06 to each page’s ideology score.<sup>10</sup>

**MEASURING PAGE HOMOGENEITY** To assess homogeneity, we use the chi-square statistic.<sup>11</sup> This statistic is used by authors such as Desmet et al. (2017) to measure overlap between ethnicity and culture and Selway (2011) to measure cross-

<sup>9</sup> We focus on ideology, though the analyses presented in this study could be done using party identification instead. In fact, page homogeneity scores based on ideology and 3-point party identification are almost perfectly correlated in our data.

<sup>10</sup> As liberals outnumber conservatives in our dataset (see Table 9.22 in the Appendix), the average ideology score across all Page Likes turns out to be less than 0.5 (0.44). As a result, a Facebook page that is liked at the same rate by liberals, moderates, and conservatives in our sample would have a page ideology of 0.44. Therefore, we add a correction factor of 0.06 to each page’s ideology score. For example, the page ideology of the Facebook page “Independent Voter” is 0.46 without the correction factor and becomes 0.52 when applying the correction. Note that this correction factor shifts the distribution of page ideology to center around 0.5 but does not affect the relative distance between different pages’ ideologies. Shi et al. (2017b) apply a similar method to adjust for uneven partisan distribution.

<sup>11</sup> As there are several other metrics could be used to measure homogeneity or audience diversity (Yamaya et al., 2020), we compare results using entropy and variance in Appendix 9.4.4.2.

cuttingness. It has the advantage that it takes into account the prior distribution of ideology in our sample when calculating homogeneity.

We consider three ideology groups ( $c$ ), i.e., *liberal*, *moderate*, and *conservative*. The fraction of likes from users with ideology  $c$  is equal to  $p_c$  (see Table 9.22). The chi-square statistic is based on comparing the distribution of ideology groups across users who have liked Facebook page  $Z$  to the distribution across users who have not liked  $Z$ . If both distributions are the same, then knowing whether a user liked Facebook page  $Z$  or not conveys no information about their ideology. If instead the distributions are distinct, then the audience of Facebook page  $Z$  is more ideologically homogeneous than the overall population.

Let  $N_c^1$  be the count of Facebook users who liked Facebook page  $Z$  and belong to ideology group  $c$  and  $N_c^0$  be the count of users in ideology group  $c$  that did not like  $Z$ . Under independence, the expected number of individuals that belong to ideology group  $c$  and like ( $i = 1$ ) or not like ( $i = 0$ ) page  $Z$  is  $N^i \times p_c$ , while the observed frequency is  $N_c^i$ . The statistic for page  $Z$  is equal to the deviation of observed values and expected values and is given by:

$$\chi(Z) = \sum_{i=0}^1 \sum_{c=0}^3 \frac{(N_c^i - N^i \times p_c)}{N^i \times p_c} \quad (6.1)$$

To ensure that this value lies between 0 and 1 we use Cramer's normalization (Equation 9.10 in Appendix 9.4.3). For one degree of freedom, a Cramer's V above 0.1 indicates a small association (the audience that likes the Facebook page is somewhat homogeneous), above 0.3 indicates medium association (homogeneous) and above 0.5 indicates a large association (very homogeneous) (Cohen, 2013).

Finally, the ideology and homogeneity scores in each category are calculated by averaging the page ideology and homogeneity scores of all pages per category, weighted by the total number of likes per page. The average homogeneity scores per Facebook category provide us with an answer to RQ1. Similarly, for each user, we average the homogeneity scores of all Facebook pages they have liked. A high homogeneity score indicates that the user tends to like more ideologically homogeneous pages. We will build a regression model to provide an observational portrait of the individual-level characteristics related to high homogeneity scores. To test H1 and answer RQ2 and RQ3, we include age, five-point ideological self-placement, and political news interest. As additional control variables, we include a mix of relevant sociodemographic variables including race, gender, family income, and educational attainment. The dependent variable is the individual homogeneity score measured by Cramer's V. Because of the nature of the dependent variable

(between 0 and 1), beta regressions are used. All replication code for this study can be found on Github.<sup>12</sup>

## 6.6 Results

The top node projection<sup>13</sup> (see Figure 6.3) reveals a clear ideological divide for the political pages based on Democratic and Republican presidential candidates. The community detection algorithm reveals distinct liberal and conservative communities with high modularity (see Table 9.29 in Appendix 9.4.5.1). For political news, there is an overlap between partisan liberal and mainstream news outlets, and they are concentrated around outlets such as The New York Times, HuffPost, and CNN. Fox News and Conservative Daily are clearly separated from the others. This seeming asymmetry recalls the analyses in Benkler et al. (2018), who describe links between conservative media and extreme partisan sites that frequently publish misleading content. For the hard news subcategory a similar pattern can be observed, except that the mainstream news outlets have a more central location in the network and form a community of their own. Strikingly, in contrast to these patterns, no ideologically distinct communities appear to emerge within the lifestyle category.

Next, we look at the page ideology distribution among liberal, moderate, and conservative respondents (see Figure 6.4). We use the overlap of these distributions as a measure of the degree to which users like ideologically similar pages (see Eady et al., 2019).<sup>14</sup> The overlap is smallest between liberals and conservatives for political pages; it increases for political news and hard news, though it is still low to moderate. The low overlapping coefficient for our news categories is especially notable in its contrast to estimates of the same statistic for news consumption through website visits on desktop, laptop and mobile devices (Guess, 2020). A likely source of this divergence is the relative absence on social media of the potentially moderating influence of news portals, aggregators, and popular mainstream website homepages — a reflection not only of differing affordances but of distinct uses and gratifications, which on Facebook may include motivations for identity signaling and affirmation in addition to simply seeking out information (Settle, 2018).

Meanwhile for lifestyle-related pages on Facebook, the distributions almost completely overlap, which again confirms our observation that lifestyle pages do not exhibit a strong ideological divide. Finally, for all types of pages, moderate participants have a larger overlap with liberals than with conservatives (see Table 9.30). In the following subsection, we analyze the ideological homogeneity of different lifestyle categories in more detail.

<sup>12</sup> [https://github.com/SPraet/facebook\\_us](https://github.com/SPraet/facebook_us)

<sup>13</sup> Results for the bottom node projection can be found in Appendix 9.4.5.1.

<sup>14</sup> We use the `overlap` package in R (Meredith and Ridout, 2014) to calculate the area lying under both of the density curves.



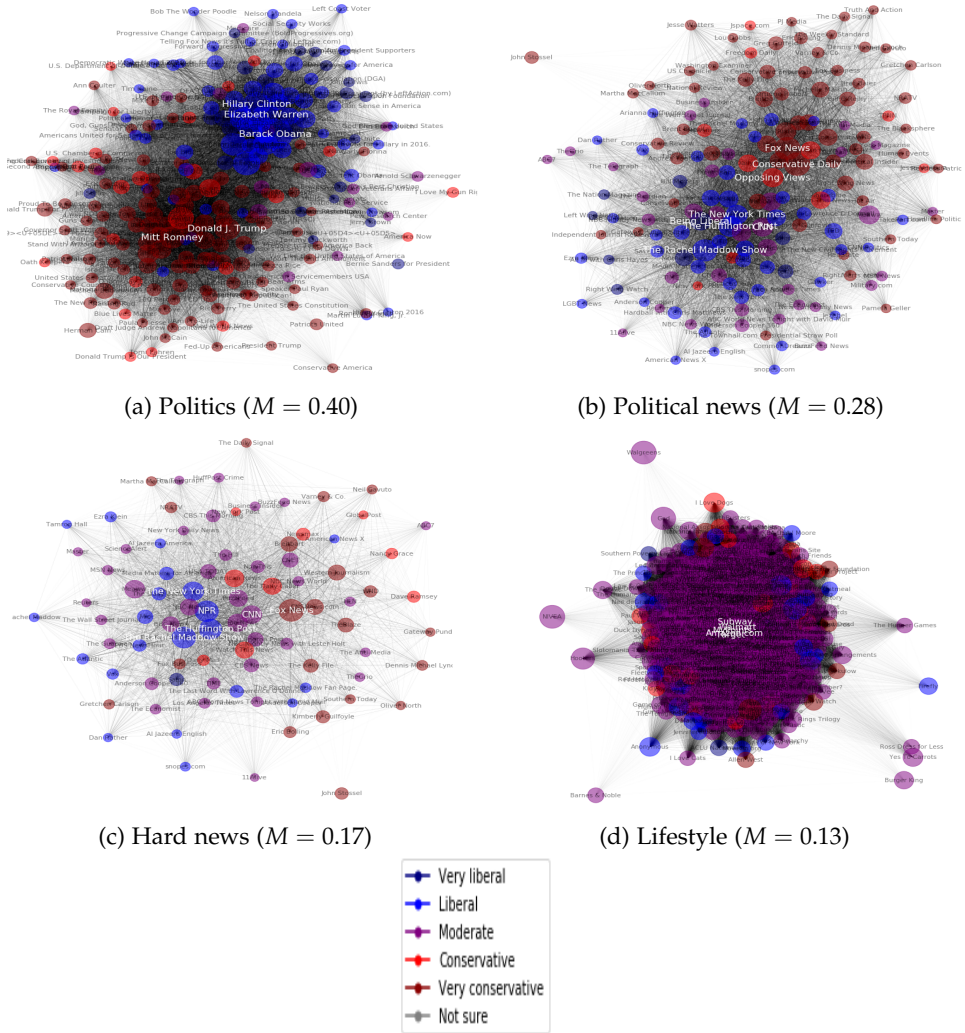


Figure 6.3: Network visualization and modularity ( $M$ ) for the top node projection of (a) political pages, (b) political news, (c) hard news, and (d) lifestyle pages. The size of the bubble represents the total number of likes of the page, and the color represents the average ideology of the audience.

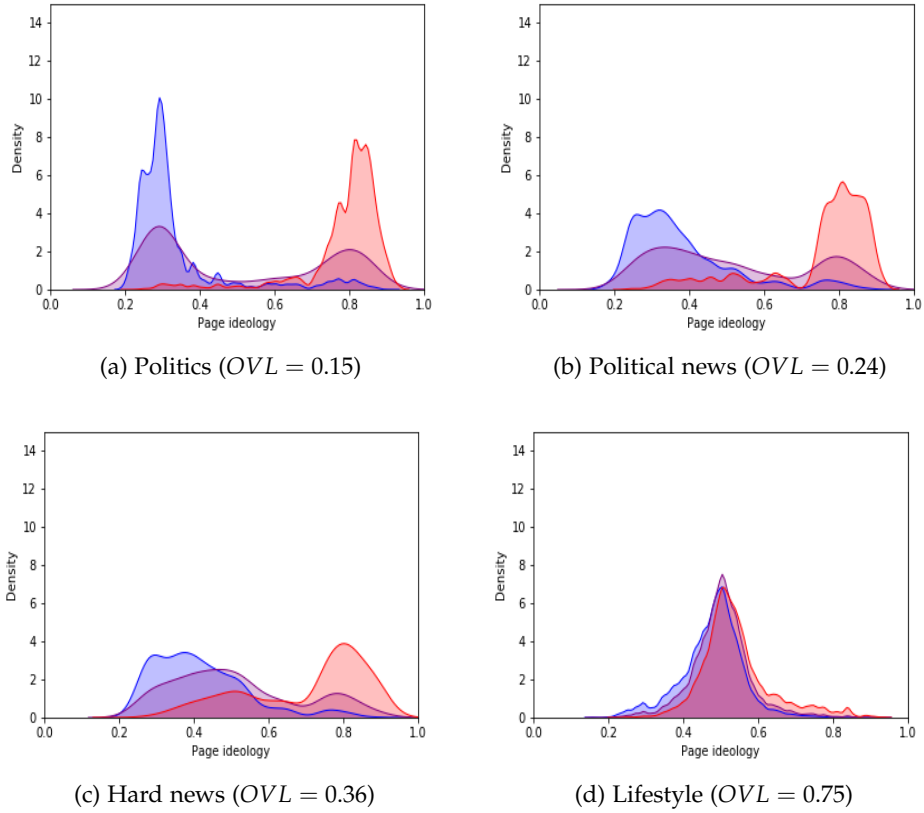


Figure 6.4: Page ideology distribution for liberals (blue), moderates (purple) and conservatives (red) when taking into account (a) political pages, (b) political news, (c) hard news, and (d) lifestyle pages; and the overlapping coefficient ( $OVL$ ) for the liberal and conservative distribution.

## 6.6.1 Facebook categories

To address RQ1, Table 6.3 summarizes the average ideology and homogeneity scores of the Facebook pages in each category. The categories **Politics** and **Political news** have the highest homogeneity scores and serve as a benchmark for polarization. Several other categories that are somewhat related to politics are relatively polarized as well: **Hard news**, **Civil society**, **Identity & religion**, **Individual opinion leaders**, and **Public figures**. Most other lifestyle categories, however, show much less polarization across ideological lines. Among the least polarized categories we find **Shopping & retail**, **Food & beverage**, and **Cars & transportation**. These pages are liked at almost equal rates by liberals, moderates, and conservatives.

Table 6.3: Weighted average (and standard deviation) for homogeneity, measured by Cramer’s V, and page ideology for all Facebook categories in the U.S., ordered from high to low homogeneity.

Category	Homogeneity	Ideology
<b>Politics</b>	<b>0.22 (0.09)</b>	<b>0.57 (0.25)</b>
<b>Political news</b>	<b>0.18 (0.10)</b>	<b>0.57 (0.23)</b>
Hard news	0.15 (0.10)	0.55 (0.20)
Civil Society	0.12 (0.09)	0.51 (0.19)
Identity & Religion	0.12 (0.05)	0.66 (0.13)
Individual opinion leaders	0.12 (0.11)	0.60 (0.17)
Public Figures	0.10 (0.09)	0.51 (0.16)
Arts & Culture	0.07 (0.06)	0.48 (0.13)
Tv Shows	0.07 (0.05)	0.48 (0.11)
Entertainment	0.07 (0.05)	0.51 (0.10)
Research & Education	0.06 (0.04)	0.46 (0.09)
Music	0.06 (0.03)	0.50 (0.10)
Interests	0.06 (0.03)	0.49 (0.10)
Movies	0.06 (0.04)	0.48 (0.10)
Sports	0.05 (0.03)	0.50 (0.08)
Services	0.05 (0.04)	0.51 (0.06)
Beauty & Health	0.04 (0.02)	0.50 (0.05)
Travel	0.04 (0.02)	0.52 (0.06)
Shopping & retail	0.04 (0.02)	0.51 (0.05)
Food & Beverage	0.04 (0.02)	0.50 (0.05)
Cars and transportation	0.04 (0.02)	0.52 (0.06)
Total	0.07 (0.07)	0.51 (0.11)

A more detailed analysis of the individual pages per category (see Appendix 9.4.5.2) sheds light on the most and least polarizing pages. For political pages (see Figure 9.7a), the Democratic and Republican political candidates hold the highest homogeneity scores (e.g.,  $V(\text{BarackObama}) = 0.39$  and  $V(\text{MittRomney}) = 0.43$ ).

Some government organizations such as NASA ( $V = 0.05$ ) and the National Park Service ( $V = 0.01$ ) are liked by a heterogeneous audience, but overall the number of pages with low homogeneity in this category is low. Next to politics, also **Political news** and **Hard news** show fairly high homogeneity scores, which is consistent with our network analysis results. Figure 9.7b and 9.7c show that Facebook audiences for news outlets are often heavily right- (Fox News [ $V = 0.38$ ,  $I = 0.77$ ] and Conservative Daily [ $V = 0.34$ ,  $I = 0.76$ ]) or left-leaning (The New York Times [ $V = 0.19$ ,  $I = 0.36$ ] and NPR [ $V = 0.28$ ,  $I = 0.32$ ]), with relatively few outlets attracting people with different ideologies (CNN [ $V = 0.06$ ,  $I = 0.45$ ], Meaww [ $V = 0.01$ ,  $I = 0.51$ ] and The Los Angeles Times [ $V = 0.00$ ,  $I = 0.51$ ]).

Even within the least-polarized categories, individual pages with high homogeneity scores do exist. Looking within the **Food & beverage** category (see Figure 9.7f), we see that, as the discussion in Settle (2018) suggests, Chick-fil-A ( $V = 0.17$ ,  $I = 0.61$ ) does have a relatively high homogeneity score in addition to its more conservative ideology rating. As Settle (2018) recounts, the chain encountered controversy in 2012 about its owner's (and charitable arm's) support for anti-gay organizations, after which activists (mainly liberals) announced a boycott of the restaurant, and others (mainly conservatives) began a counter-boycott. In this way, Chick-fil-A became a politicized topic such that, apparently, by the time of our data collection in 2016, liking the Facebook page of the fast-food chain could be seen as an endorsement of the political views of the company. In the opposite sense, the ice-cream brand Ben & Jerry's openly promotes progressive values and expresses support for social and environmental justice initiatives around the country. Though homogeneity is low ( $V = 0.09$ ,  $I = 0.36$ ), it is relatively high compared to other pages in the food category, and the brand is predominantly liked by liberal users.

Likewise, **Sports** (see Figure 9.7e) can also become caught in the partisan currents of the larger society as a result of symbolic actions and outspoken statements of its players. For example, Tim Tebow ( $V = 0.11$ ,  $I = 0.71$ ) has a predominantly conservative following, and some NASCAR ( $V = 0.11$ ,  $I = 0.57$ ) drivers have publicly supported Republican candidates. The Olympics ( $V = 0.09$ ,  $I = 0.33$ ) and the Pittsburgh Steelers ( $V = 0.10$ ,  $I = 0.43$ ) have a more liberal audience, while Serena Williams ( $V = 0.13$ ,  $I = 0.47$ ) fans are more moderate on average. Still, the majority of sports pages appear to unite people with different ideologies, and have very low homogeneity scores including the Boston Red Sox ( $V = 0.00$ ,  $I = 0.48$ ), New England Patriots ( $V = 0.02$ ,  $I = 0.48$ ), and New York Yankees ( $V = 0.02$ ,  $I = 0.51$ ). In 2016, LeBron James ( $V = 0.04$ ,  $I = 0.43$ ) was popular across the ideological spectrum,<sup>15</sup> though we suspect that this might have shifted in the period after our data collection given his subsequent criticisms of President Trump.

Our results show both similarities and contrasts with Shi et al. (2017b), who study partisan divisions in the U.S. by analyzing Twitter co-following networks. They

<sup>15</sup> Note that even though the average ideology of LeBron James is equal to that of the Pittsburgh Steelers, the homogeneity score of the first is much lower and thus his audience is more diverse.

too find cultural dimensions other than religion to be substantially less polarized than political domains. A striking difference from our study, however, is that they find news and media among the dimensions that cut across the political divide. Such interpretations illustrate the difficulty of establishing empirical benchmarks for polarization in addition to comparing estimates of magnitude across different measures.

### 6.6.2 User-level analysis

We now turn to the individual level. On average, conservative users like more pages on Facebook, and a slightly higher percentage of the pages they like are political (Table 6.4). Similar to Eady et al. (2019), we also look at the proportion of liberals and conservatives whose Page Likes include at least 5% of pages at the right and left ends of the spectrum, respectively. For each group of pages, we consider “left-leaning” pages as pages with a page ideology score that is lower than the 70th percentile of all pages liked by liberal participants, and “right-leaning” pages are pages with a page ideology score higher than the 30th percentile of all pages liked by conservative participants. Examples of pages at these percentiles of political, news, and lifestyle pages can be found in Table 6.5.

Table 6.4: Average number of Page Likes per ideology

	Politics	Political News	Hard News	Lifestyle	All pages
Liberals	21.79 (8%)	12.45 (5%)	7.37 (3%)	234.46 (85%)	276.07
Moderates	17.83 (6%)	10.66 (4%)	7.03 (2%)	252.03 (88%)	287.55
Conservatives	41.62 (13%)	22.64 (7%)	11.51 (3%)	253.32 (77%)	329.09

Table 6.5: Examples and ideology score of 70th percentile of all pages liked by liberals and 30th percentile of all pages liked by conservatives per group of pages.

	70th percentile liberal Page Likes	30th percentile conservative Page Likes
Politics	Chelsea Clinton (0.32)	Mitt Romney (0.78)
Political news	BBC News (0.40)	Conservative Daily (0.76)
Hard news	Washington Post (0.46)	Fox 5 New York (0.62)
Lifestyle	NIVEA (0.52)	Amazon (0.49)

We find that 9% of liberals have at least 5% of pages to the right of Mitt Romney among their political Page Likes (see Table 6.6). Similarly, the political Page Likes of 8% of conservatives consist of at least 5% pages to the left of Chelsea Clinton. For news pages, conservatives are more likely to follow “left-leaning” pages than the other way around, a finding that corresponds to the results of Eady et al. (2019).

For lifestyle pages, almost all liberals and conservatives follow at least 5% opposite-leaning pages.

Table 6.6: Proportion of liberals with at least 5% right-leaning Page Likes (page ideology higher than 30th percentile of all pages liked by conservatives) and of conservatives with at least 5% left-leaning Page Likes (page ideology lower than 70th percentile of all pages liked by liberals) per group of pages.

	Liberals - right Likes	Conservatives - left Likes
Politics	9%	8%
Political news	7%	19%***
Hard news	15%	28%***
Lifestyle	97%	97%

\* $p < .1$ ; \*\* $p < .05$ ; \*\*\* $p < .01$

Two-tailed Z-test for two population proportions

It is possible that politically engaged individuals consciously choose to like political or news pages that have opposing views in order to stay informed about “the other side.” To test this, we use a two-tailed  $t$ -test to compare the average number of political Page Likes by liberals and conservatives that like opposing pages, to the average number of political Page Likes by liberals and conservatives overall (Table 6.7). Liberals that like news pages with opposing views have more political Page Likes on average,<sup>16</sup> and this is also the case for conservatives when we consider hard news pages. In general, more politically engaged individuals are thus more likely to follow news pages that contain opposing views, but not more likely to follow opposing political candidates.

In Figure 6.5, we zoom in on the individual homogeneity scores of our participants. For political and news pages, conservatives have slightly higher homogeneity scores than liberals and moderates. For lifestyle pages, the majority of homogeneity scores are low (below 0.1) regardless of ideology.

To gain a more complete understanding of individual-level determinants of Page Like homogeneity, we perform a beta regression that includes sociodemographic variables in addition to ideological self-placement and political news interest as predictors. Total number of Page Likes is also included as a control variable to account for individual-level differences in engagement with the platform.

The results for the four page categories are shown in Table 6.8. We find that conservatives (whether strong or not) are more likely to like homogeneous pages regardless of category (RQ2). For liberals this is true for those who are the strongest

<sup>16</sup> We do not find significant results for political and lifestyle pages.

Table 6.7: Average number of political Page Likes for liberals with at least 5% right-leaning Page Likes and for conservatives with at least 5% left-leaning Page Likes per group of pages.

	Liberals - right Likes	Conservatives - left Likes
Politics	26.23	23.77
Political news	38.11***	57.64
Hard news	42.68***	59.25**
Lifestyle	22.14	41.18

\*  $p < .1$ ; \*\*  $p < .05$ ; \*\*\*  $p < .01$

Two-tailed T-test for the means of two independent samples

Mean political Page Likes for liberals is 21.79

Mean political Page Likes for conservatives is 41.62




Figure 6.5: Homogeneity distribution for liberals (blue), moderates (purple) and conservatives (red) when taking into account (a) political pages, (b) political news, (c) hard news, and (d) lifestyle pages.

Table 6.8: Determinants of individual homogeneity per category

	Politics	Political news	Hardnews	Lifestyle
Age: 30-44	-0.094 (0.058)	0.116 (0.085)	-0.035 (0.093)	-0.012 (0.045)
Age: 45-65	-0.060 (0.055)	0.149* (0.080)	0.018 (0.087)	0.034 (0.043)
Age: Over 65	-0.002 (0.062)	0.283*** (0.087)	0.202** (0.095)	0.265*** (0.048)
Black	0.007 (0.055)	-0.347*** (0.076)	-0.360*** (0.086)	-0.141*** (0.042)
Hispanic	-0.096 (0.067)	-0.175** (0.089)	-0.265*** (0.098)	0.038 (0.050)
Other Race	-0.169** (0.066)	-0.355*** (0.086)	-0.427*** (0.098)	-0.107** (0.049)
Female	0.008 (0.032)	-0.024 (0.042)	-0.071 (0.047)	-0.117*** (0.025)
Income	0.007* (0.004)	-0.005 (0.005)	0.003 (0.006)	0.003 (0.003)
Education	-0.019 (0.011)	0.013 (0.015)	0.027 (0.017)	0.021** (0.009)
Very Liberal	0.087* (0.048)	0.047 (0.064)	0.175** (0.072)	0.168*** (0.037)
Liberal	0.021 (0.046)	-0.039 (0.062)	0.061 (0.069)	0.074** (0.036)
Conservative	0.208*** (0.044)	0.278*** (0.057)	0.302*** (0.066)	0.154*** (0.034)
Very Conservative	0.158*** (0.058)	0.491*** (0.069)	0.549*** (0.079)	0.200*** (0.044)
Political news interest	0.092*** (0.023)	0.181*** (0.031)	0.190*** (0.034)	0.132*** (0.017)
Number of likes	-0.00002** (0.00001)	-0.00002* (0.00001)	-0.00001 (0.00001)	-0.00003*** (0.00001)
Constant	-1.055*** (0.087)	-1.556*** (0.120)	-1.689*** (0.133)	-2.394*** (0.067)
N	826	774	740	1,085
Pseudo R <sup>2</sup>	0.074	0.189	0.195	0.273

\*p &lt; .1; \*\*p &lt; .05; \*\*\*p &lt; .01

Beta regressions with survey weights applied. Reference categories are Age: 18-29, White race, Male gender and Moderate ideology.

Income ranges from 1 to 31, Education from 1 to 6, and political news interest from 1 to 4.



liberals, except for the lifestyle category, where liberals (whether strong or not) are also more likely to like homogeneous pages.<sup>17</sup> Higher political news interest is also associated with higher individual homogeneity scores across all categories.

These results suggest that polarization of Page Likes in “non-political” domains remains limited to more politically active individuals, which confirms Hypothesis 1. However, when we include party identification instead of ideology (see Table 9.27), we find Republicans to be more likely to like homogeneous pages, though partisanship strength does not seem to matter. For lifestyle pages, we find no relationship with party identification. Our results thus suggest that polarization in lifestyle pages is predominantly tied to ideological strength and political interest rather than strong partisanship.

Turning to RQ3, we find that older age is predictive of greater homogeneity in Page Likes, but only for news and lifestyle pages. Our results for news pages are consistent with Guess et al. (2019b), who find — using the same underlying data source as this analysis — that conservatives and people over the age of 65 were more likely to share “fake news” on Facebook in 2016, all else equal. This suggests that page liking patterns may be part of a process in which online misinformation reaches people’s social media feeds, thereby increasing the likelihood of engaging with it and sharing it with one’s social connections (e.g., Grinberg et al., 2019).

The analysis in Table 6.8 reveals other relationships worth exploring in future research. Individuals with higher educational attainment are more likely to like more homogeneous lifestyle pages but not more likely to like homogeneous political or news pages, a pattern consistent with research finding that highly educated people are more likely to make consumer decisions that reflect their political leanings (Newman and Bartels, 2011). Similarly, gender only has a significant effect within lifestyle pages, but we speculate that the effect could vary depending on the lifestyle subcategories.

Liking explicit political content on Facebook is a form of political participation or endorsement, while liking lifestyle pages may seem apolitical at first sight. Our findings suggest that the characteristics of individuals who exhibit high levels of ideological homogeneity are different for explicitly and implicitly political pages. While most research on polarization and echo-chamber dynamics has focused on networks around explicit political content, an analysis of these lifestyle categories reveals a subtler form of political homophily.

---

<sup>17</sup> Moderate is the reference category for ideology.

## 6.7 Conclusion

In light of increasing discussions about political divides in the United States, we explore polarization in political and traditionally “non-political” domains on social media. Our results from analyzing Facebook Like data suggest that ideological divides are large in relatively political domains such as news and media, civil society, and religion but much less pronounced in areas such as culture, food, and sports. Our findings show that polarization does not permeate society as a whole: Lifestyle endeavors still offer cross-cutting spaces, and polarization, when it does emerge, seems limited to a narrow set of politicized examples. Considering that Facebook users primarily engage with non-political Facebook Likes, our findings add nuance to debates about the divisive nature of social platforms.

At the individual level, we find that polarization in page liking patterns is more associated with politically active individuals. If political polarization were thoroughly permeating society, we think that we would not have observed a higher likelihood among respondents with stronger ideological preferences and political interest of endorsing more homogeneous lifestyle pages. Furthermore, when we distinguish between polarization in political and non-political domains, we find that individuals who exhibit high levels of political homophily are different. For example, highly educated people are more likely to make lifestyle choices that are related to their political views, while they are less likely to be in echo chambers reflecting the political pages that they follow. This finding has potentially important implications for research on online polarization.

Given our findings, then, why do narratives of enduring political divides in non-political domains persist? One explanation is that people draw inferences on the basis of vivid but unrepresentative examples, as our analysis of Chick-fil-A and prominent sports figures suggests.<sup>18</sup> Similarly, people have exaggerated perceptions of the differences between the parties, both in terms of demographic composition and lifestyle tendencies (Ahler and Sood, 2018).<sup>19</sup> Social media itself may drive these misperceptions by fueling cycles of engagement with content that promotes disparagement of partisan outgroups (e.g., Barberá et al., 2015). Future research should consider how users’ online social endorsements interact with these dynamics over time, especially as a possible window into the politicization of figures and brands. As Settle (2018) illustrates, the process by which one’s political preferences come to influence seemingly distinct consumer and lifestyle choices can emerge unexpectedly as a result of both elite actions and mass mobilization. Even though we show these cases to be the exception, they demonstrate how the coexistence of

<sup>18</sup> This is also borne out in polling, which tends to emphasize these vivid examples in addition to brands known to be polarizing, such as media organizations. See <https://morningconsult.com/polarizing-brands-2018/>.

<sup>19</sup> The New York Times’ recent feature asking readers to guess people’s vote preferences from the contents of their refrigerators illustrates the limited predictive value of partisan stereotypes. See <https://www.nytimes.com/interactive/2020/10/27/upshot/biden-trump-poll-quiz.html>.

political and other identities on social media leaves users vulnerable to mechanisms of social polarization.

Social media data offer a rich source of information about individuals' revealed social and lifestyle preferences, at a resolution that would be difficult to attain with traditional survey techniques. At the same time, the collection of online behavioral data comes with its own set of ethical and privacy challenges (Stier et al., 2019). Drawing inferences from online data should be performed with caution since ignoring offline behavior may leave us with a distorted view. Still, linking digital trace data with survey data helps us to understand the relationship between lifestyle preferences and politics and to map the landscape of political culture — both its fault lines and its areas of overlap.



## Polarization in Belgium compared to the U.S.

---

The previous two chapters demonstrate how Facebook Likes offer a rich source of information about individuals' revealed social and lifestyle preferences. Combined with survey information, these digital traces provide us with a unique opportunity to study (online) polarization. In this final chapter, we compare ideological disunity in political and non-political domains in Belgium (a multi-party political system), to the U.S. (a two-party system). We find that polarization in the categories politics and news and media is noticeably lower in Belgium and that less individuals in Belgium are being exposed to large amounts of highly homogeneous content. Individuals that exhibit higher levels of individual homogeneity are often more politically active. In contrast to the U.S., political affiliation is more likely to seep through in lifestyle choices for individuals with a left ideological leaning. But, similar to the U.S., we find the majority of lifestyle pages to unite people with different political views.

## 7.1 Introduction

Ever more countries are facing increasing levels of (affective) polarization (Iyengar et al., 2012; Lelkes, 2016; Westwood et al., 2018). Yet, the main body of research focusing on online and offline polarization has been applied to the American two-party system (Wagner, 2020). Urman (2020) shows that the levels of polarization on social media platforms can vary depending on countries' electoral rules and party systems. Belgium has a fragmented multi-party system with large coalition governments, while the U.S. is a two-party majoritarian system. So how would results for the Belgian case differ from what we found in the previous chapter? Westwood et al. (2018) compare partyism in four countries and find that both in the U.S. and in Belgium, citizens are more trusting of co-partisans and less trusting of opposing partisans. In Belgium, partisan animus is conditioned by ideological proximity, i.e. partisans are more distrusting of parties furthest from them in the ideological space. While American citizens frequently display their political affinity (e.g. bumper stickers or campaign merchandise), citizens' political affiliations are often less discernible in Belgium, which weakens the divisive impact of partyism. Compared to the U.S., we expect less homogeneity in political Facebook pages in Belgium, due to the presence of center parties. Furthermore, news brands traditionally enjoy high levels of trust in Belgium, especially in Flanders (Newman et al., 2019b), and we expect them to be hardly politically polarized compared to the U.S. Lastly, although lifestyle pages, especially in the categories civil society and culture, appear to be predictive for ideology and political preference (see Chapter 5), we do not expect the majority of lifestyle pages to be polarized to worrisome extent. In contrast, we trust to find plenty of evidence of cross-cutting behavior in Belgium, in line with our nuanced findings for American society.

## 7.2 Data and methods

The data collection process for Belgium and the U.S. is described in Chapter 5 and Chapter 6 respectively. Both datasets combine survey data with Facebook Like data and are highly similar. However, a few differences complicate direct comparison of the results. First, the Facebook categories in both datasets are compiled differently. In the Belgian data, we make use of the page categories indicated on the Facebook pages and combine them into 20 categories. The categories in the U.S. data are based on manual labeling and slightly different categories. We will only consider comparable categories in a discussion of the results. Also, we will no longer make a distinction between political news and hard news. Second, in Belgium, self-reported ideology was indicated on a scale from 0 to 10, in the U.S. from 1 to 5. In Belgium we'll consider the ideology scores 0 and 1 as "very left"; 2 and 3 as "left"; 4, 5 and 6 as "center"; 7 and 8 as "right"; and higher than 8 as "very right". We combine the scores in three ideology classes "left", "center", and "right". Finally, the amount, region, and time period of data collection differs. We have collected data of more than 6,500 participants in one region in Belgium (Flanders) in 2018, compared to

around 1,200 respondents across all states of the U.S. However, despite the large difference in geographical region covered, both regions can be considered to be integrated societies, in terms of common language and culture.

We apply the same methods described in Chapter 6 to the Belgian Facebook Likes and compare the results the U.S. We calculate page ideology ( $I$ ) and homogeneity ( $V$ ) and average this over categories and users to answer two main questions: How ideologically homogeneous are political and lifestyle categories on Facebook? (RQ1) and which characteristics are associated with a greater likelihood of “liking” ideologically homogeneous pages? (RQ2)

### 7.3 Results


We first compare the page ideology distribution among conservative/left, moderate/center and conservative/right leaning respondents (see Figure 7.1). With regard to political pages, we immediately notice the presence of center political pages in Belgium. The density is high around 0.5, especially among center voters. A political center is absent in the U.S., since there are only two dominant parties. Since both left and right respondents in Belgium like pages with a center page ideology, the overlap between both distributions is much higher than in the U.S. For news pages, the distributions almost completely overlap in Belgium, while there is a large divergence between liberals and conservatives in the U.S. In Belgium, news sources seem to be trusted (or at least “liked” on Facebook) across the ideological spectrum. Finally, for lifestyle-related pages on Facebook, oddly, we find the overlap between distributions in Belgium to be slightly lower than for news pages, and also lower than in the U.S. (see Table 9.31). In the following subsection, we analyze the ideological homogeneity of different lifestyle categories in more detail.

#### 7.3.1 Facebook categories


Table 7.1 summarizes the average homogeneity scores of the Facebook pages in each category.<sup>1</sup> The results for Belgium largely correspond to the predictive power of the categories we analyzed in Chapter 5. In both countries, the category **Politics** has of course the highest homogeneity score, although it is remarkably lower in Belgium. In the U.S., a few other categories are relatively polarized as well: **News**, **Civil society**, and **Public figures**, which are hardly polarized in Belgium. All other lifestyle categories show little to no polarization across ideological lines.

A more detailed analysis of the individual pages per category sheds light on the most and least polarizing pages. Similar to the U.S., for political pages (see Figure 7.2b),


<sup>1</sup> Remember from Section 6.5 that a Cramer’s  $V$  above 0.1 indicates a small association (the audience that likes the Facebook page is somewhat homogeneous), above 0.3 indicates medium association (homogeneous) and above 0.5 indicates a large association (very homogeneous)




(a) Politics U.S. ( $OVL = 0.15$ )




(b) Politics Belgium ( $OVL = 0.36$ )




(c) News U.S. ( $OVL = 0.36$ )



(d) News Belgium ( $OVL = 0.72$ )



(e) Lifestyle U.S. ( $OVL = 0.75$ )



(f) Lifestyle Belgium ( $OVL = 0.68$ )

Figure 7.1: Page ideology distribution for liberal/left (blue), moderate/center (purple) and conservative/right (red) respondents when taking into account political pages, news pages, and lifestyle pages; and the overlapping coefficient ( $OVL$ ) for the left and right distribution.



Table 7.1: Weighted average (and standard deviation) for homogeneity, measured by Cramer's  $V$ , for the Facebook categories in the U.S. and Belgium.

Category	U.S.	Belgium
<b>Politics</b>	<b>0.22 (0.09)</b>	<b>0.09 (0.08)</b>
News	<b>0.15 (0.10)</b>	0.06 (0.05)
Civil Society	<b>0.12 (0.09)</b>	0.06 (0.05)
Public Figures	<b>0.10 (0.09)</b>	0.04 (0.03)
Arts & Culture	0.07 (0.06)	0.05 (0.03)
Tv Shows	0.07 (0.05)	0.03 (0.02)
Entertainment	0.07 (0.05)	0.04 (0.03)
Research & Education	0.06 (0.04)	0.03 (0.02)
Music	0.06 (0.03)	0.04 (0.02)
Movies	0.06 (0.04)	0.03 (0.02)
Sports	0.05 (0.03)	0.04 (0.03)
Services	0.05 (0.04)	0.05 (0.03)
Beauty & Health	0.04 (0.02)	0.04 (0.03)
Travel	0.04 (0.02)	0.03 (0.01)
Shopping & retail	0.04 (0.02)	0.04 (0.03)
Food & Beverage	0.04 (0.02)	0.03 (0.02)
Total	0.07 (0.07)	0.04 (0.04)

left and right political parties and candidates hold the highest homogeneity scores (e.g., Groen ( $I = 0.55$ ,  $V = 0.29$ ) and NVA ( $I = 0.55$ ,  $V = 0.29$ )). In contrast however, there are large pages located towards the center. The Facebook page of the Christian Democratic party CD&V ( $I = 0.55$ ,  $V = 0.04$ ) is not only liked by center voters but also left and right voters. Remarkably, despite the fact that they can't vote for him — or maybe it is because of that — Barack Obama ( $I = 0.53$ ,  $V = 0.00$ ) is by far the most popular politician in Belgium and he can persuade both the political left and right.

With regard to news and media pages, Figure 7.2d shows that besides some heavily left or right leaning news sources (e.g., De Wereld Morgen ( $I = 0.53$ ,  $V = 0.00$ ) and SCEPTR ( $I = 0.78$ ,  $V = 0.22$ )), the majority of news pages is concentrated in the center with relatively low homogeneity scores. The public broadcaster, VRT ( $I = 0.52$ ,  $V = 0.03$ ) attracts a balanced audience of left, center and right voters. Indeed, recent research confirms that the VRT succeeds in providing balanced and impartial news to the public (Raats et al., 2021).<sup>2</sup>

All lifestyle categories are predominantly heterogeneous, although individual pages with high homogeneity scores do exist. Within the **Food & beverages** category (see Figure 7.2f), EVA ( $I = 0.33$ ,  $V = 0.17$ ), an organization that promotes vegetarian alternatives, does have a relatively high homogeneity score in addition to its left

<sup>2</sup> They examined the impartiality of the public broadcaster based on an analysis of their content, production and public perception.

ideology rating. Plant-based food is put forward as an important part of the solution to our environmental issues, and as we also found in Chapter 3, the environment can be considered a predominantly left theme. It is therefore little surprising that McDonald's ( $I = 0.65$ ,  $V = 0.08$ ) is situated at the other end of the spectrum. Besides these few examples, ideological homogeneity in the food category is in fact limited. And whatever our dietary or political preferences, we all shop food at Albert Heijn ( $I = 0.53$ ,  $V = 0.00$ ).

### 7.3.2 User-level analysis


In Figure 7.3, we zoom in on the individual homogeneity scores of our participants. For political pages, right voters have slightly higher homogeneity scores than left and especially center voters. For news and lifestyle pages, the majority of homogeneity scores is low (below 0.1) regardless of ideology. This is in sharp contrast to the U.S., where the proportion of participants with higher homogeneity scores is much higher, especially for political pages but also for the other categories (see Table 7.2). We can conclude that less individuals in Belgium are exposed to highly homogeneous content.

Table 7.2: Proportion of respondents with individual homogeneity score higher than 0.1


	U.S.	Belgium
Politics	0.71	0.23
News	0.44	0.06
Lifestyle	0.18	0.00

To gain a more complete understanding of individual-level determinants of Page Like homogeneity, we perform a beta regression (see Table 7.3) that includes sociodemographic variables in addition to ideological self-placement and political news interest as predictors. Total number of Page Likes is also included as a control variable to account for individual-level differences in engagement with the platform.

Left and right voters are more likely than center voters to like homogeneous pages regardless of category. The effect is larger for stronger ideologies. For political pages, especially (very) right voters are more likely to like homogeneous pages, a finding that is confirmed by the homogeneity distribution for right voters in Figure 7.3b. Conversely, for the news and lifestyle category, left voters are more likely to like homogeneous pages (again confirmed by Figure 7.3f). This is in contrast to the U.S., where for all categories, conservatives are more likely to like homogeneous content. Higher political news interest is also associated with higher individual homogeneity scores across all categories. These results suggest that polarization of Page Likes in “non-political” domains remains limited to more politically active individuals, and is




(a) Politics U.S.




(b) Politics Belgium

Figure 7.2: Ideology and homogeneity scores for the Facebook pages per category. The magnitude of the circle represents the total number of likes of the Facebook page.




(c) News U.S.




(d) News Belgium

Figure 7.2: (Continued) Ideology and homogeneity scores for the Facebook pages per category. The magnitude of the circle represents the total number of likes of the Facebook page.




(e) Food &amp; Beverages U.S.




(f) Food &amp; Beverages Belgium

Figure 7.2: (Continued) Ideology and homogeneity scores for the Facebook pages per category. The magnitude of the circle represents the total number of likes of the Facebook page.


# POLARIZATION IN BELGIUM COMPARED TO THE U.S.




(a) Politics U.S.




(b) Politics Belgium




(c) News U.S.



(d) News Belgium



(e) Lifestyle U.S.



(f) Lifestyle Belgium

Figure 7.3: Average homogeneity distribution for liberal/left (blue), moderate/center (purple) and conservative/right (red) respondents when taking into account political pages, news and media pages, and lifestyle pages.

Table 7.3: Determinants of individual homogeneity with ideology

	All pages (1)	Political (2)	News (3)	Lifestyle (4)
Age: 25-55	0.079*** (0.009)	0.063** (0.026)	0.047*** (0.013)	0.057*** (0.008)
Age: 55+	0.175*** (0.014)	0.351*** (0.039)	0.035* (0.021)	0.048*** (0.013)
Female	-0.048*** (0.009)	-0.061** (0.026)	-0.036*** (0.013)	-0.044*** (0.008)
Education	0.006* (0.003)	-0.052*** (0.009)	0.010** (0.005)	0.017*** (0.003)
Very Left	0.204*** (0.014)	0.358*** (0.040)	0.209*** (0.021)	0.216*** (0.012)
Left	0.118*** (0.010)	0.148*** (0.031)	0.119*** (0.015)	0.129*** (0.009)
Right	0.059*** (0.011)	0.268*** (0.031)	0.032** (0.016)	0.010 (0.010)
Very Right	0.211*** (0.017)	0.671*** (0.046)	0.111*** (0.027)	0.098*** (0.016)
Political news interest	0.071*** (0.004)	0.122*** (0.012)	0.052*** (0.006)	0.044*** (0.004)
Number of likes	0.00000 (0.00001)	-0.00001 (0.00002)	0.00001 (0.00001)	0.00001 (0.00001)
Constant	-3.396*** (0.019)	-2.719*** (0.057)	-3.030*** (0.028)	-3.462*** (0.017)
N	5,573	3,981	5,186	5,557
Pseudo R <sup>2</sup>	0.176	0.139	0.055	0.136

\*p < .1; \*\*p < .05; \*\*\*p < .01

Beta regressions. Reference categories are Age: 25-, Male gender, and Center ideology  
Education ranges from 1 to 7, and political news interest from 1 to 5.

more tied to the political left. We reach a similar conclusion when we include party preference instead of ideology (see Table 7.4), where we find left wing parties to be stronger predictors for homogeneity in news and lifestyle pages, and right wing parties for political pages.

Table 7.4: Determinants of individual homogeneity with party preference

	All pages (1)	Political (2)	News (3)	Lifestyle (4)
Age: 25-55	0.074*** (0.009)	0.059** (0.026)	0.041*** (0.013)	0.054*** (0.008)
Age: 55+	0.175*** (0.013)	0.347*** (0.038)	0.022 (0.021)	0.051*** (0.013)
Female	-0.048*** (0.009)	-0.050** (0.025)	-0.031** (0.013)	-0.045*** (0.008)
Education	0.008** (0.003)	-0.048*** (0.009)	0.012*** (0.005)	0.018*** (0.003)
PVDA	0.186*** (0.018)	0.466*** (0.053)	0.177*** (0.026)	0.176*** (0.016)
Groen	0.150*** (0.015)	0.340*** (0.047)	0.121*** (0.022)	0.150*** (0.014)
Sp.a	0.099*** (0.018)	0.141** (0.055)	0.117*** (0.026)	0.101*** (0.016)
Open VLD	0.019 (0.017)	0.178*** (0.051)	0.012 (0.024)	0.008 (0.015)
NVA	0.188*** (0.016)	0.738*** (0.048)	0.098*** (0.024)	0.071*** (0.015)
Vlaams Belang	0.138*** (0.028)	0.587*** (0.078)	0.101** (0.042)	0.076*** (0.026)
Political news interest	0.082*** (0.004)	0.148*** (0.011)	0.060*** (0.006)	0.053*** (0.004)
Number of likes	0.00000 (0.00001)	-0.00001 (0.00002)	0.00002 (0.00001)	0.00001 (0.00001)
Constant	-3.475*** (0.022)	-2.993*** (0.066)	-3.088*** (0.033)	-3.521*** (0.020)
N	5,739	4,097	5,343	5,723
Pseudo R <sup>2</sup>	0.173	0.173	0.045	0.119

\*p < .1; \*\*p < .05; \*\*\*p < .01

Beta regressions. Reference categories are Age: 25-, Male gender, and CD&V  
Education ranges from 1 to 7, and political news interest from 1 to 5.

Furthermore, male gender and older age are predictive of greater homogeneity in Page Likes. However, in contrast with our findings in the U.S., the oldest age category is not more likely to like more homogeneous news pages. Since Belgians traditionally have had high trust in mainstream media, the older generation may also fall back on these mainstream sources on social platforms. In contrast, Facebook users between 25 and 55 are possibly exploring more alternative news sources. Finally, individuals with higher educational attainment are more likely to like more



homogeneous news and lifestyle pages but not more likely to like homogeneous political pages. Once again, we find different results for explicitly and implicitly political pages—where we consider news pages in Belgium as implicitly political pages since the average homogeneity in this category is very low.

## 7.4 Conclusion

In light of discussions about increasing political divides in countries worldwide, we need more research outside of the United States that concern polarization on social media. Using Facebook Like data in Belgium (multi-party) and the U.S. (two-party), we explore differences in polarization in political and traditionally “non-political” domains on social media. Our results suggest that, as expected, ideological divides in political domains are much less outspoken in Belgium compared to the U.S. This is arguably due to the presence of a political center that unites voters with different ideologies. Indeed, Urman (2020) also find polarization on social media to be higher in two-party compared to multi-party systems. With regard to lifestyle categories, we do not find evidence of strong polarization in neither of the countries. We should note that the results and conclusions in this study are based on relatively popular Facebook pages (a minimum of 30 likes in our dataset) and that we cannot analyze polarization on smaller pages due to statistical limitations.

At the individual level, we can conclude that less individuals in Belgium are at the danger of being exposed to highly homogeneous content only. Individuals that do exhibit higher levels of individual homogeneity are often more politically active individuals (with stronger ideological preferences and higher political news interest). Especially the ideological right tends to like homogeneous political pages. However, we find left voters to be more likely to like homogeneous pages for the news and lifestyle category, which is in contrast to our findings in the U.S. Furthermore, we do not find the oldest age category to be related to liking homogeneous news pages, which was an important predictor in the U.S. and could also be related to the likelihood of sharing fake news. This could possibly be explained by high mainstream media trust among the older population in Belgium. Finally, the Belgian data confirms that individuals who exhibit high levels of individual homogeneity are different for explicit and implicit political domains. As argued before, this finding has potentially important implications for research on online polarization.

To conclude, although we found Facebook pages to be highly predictive for ideology and political leaning in Chapter 5, we do not find evidence of large political polarization in general. The majority of lifestyle pages unites people with different ideological views, and polarization seems limited to a narrow set of politicized examples. When combined, (non-political) Facebook likes are very informative of our underlying ideological beliefs, and they are very powerful in distinguishing left from right voters. Yet, based on the results of this analysis, we can—fortunately—conclude that what unites us is still more prominent than what divides us.



## Part III

# CONCLUSIONS



## Conclusions

---

*“Technology is a very human activity - and so is the history of technology.”*

*Kranzberg’s Sixth Law of Technology*

The rise of social media has not only reshaped our social, economic, and political lives, it has also led to a data revolution. This thesis explored the opportunities of data mining, text mining, and network analysis techniques to process large amounts of observational data and learn about political behavior and communication on a fine-grained level. Five unique studies show that social media data provide a unique and rich source of information, but also come with ethical, technical and methodological challenges. This final chapter summarizes the main empirical and methodological contributions from this work and presents avenues for further research. I suggest recommendations to enable more impactful research in the field of computational social science with the use of social media data. Finally, I end with an outlook on the future role of social media in our society, and how humans can interact more consciously with technology to build a better online future.

## 8.1 Main findings

Social media has shifted who controls, consumes and distributes political information. A first vital change is that it allows politicians and political parties to shape their content and communicate directly with voters without the intervention of traditional media. **Part I** studies online communication by political actors on Twitter. Twitter is probably the most widely accessible form of party communication today, characterized by high individual autonomy, temporal adaptability, and interaction potential. Political parties use Twitter as a way to communicate messages to the media and the public, like a press release, even in countries with low or elite-only Twitter penetration. Next to that, Twitter data is also easily accessible for researchers, which has made it a popular source to study political communication and reassess classical communication theories in the digital age.

In **Chapter 3** we contribute methodologically to the increasing scholarly interest in parties' issue communication and strategies on social media. Traditionally, this field relied on manual coding or dictionary-based approaches to analyze the frequency of issue communication in party manifestos, campaign ads and press releases. Nowadays, the sheer volume of available texts can no longer be processed with manual encoding. Additionally, the volatile and less formal nature of social media text complicate the accurate functioning of dictionaries and other issue classification methods on this type of data. Therefore, we propose a data-driven exploratory approach to analyze issue communication at the party level, without the need to classify texts into predefined categories upfront. We systematically compare the value and shortcomings of three text representation approaches: (1) an expert-driven approach based on dictionaries, (2) a data-driven approach based on a bag of words method, and (3) another data-driven approach based on topic modeling. Our methods are applied to two years of Twitter data from six Belgian political parties to analyze which issues separate the online communication of one party from that of the other parties, and how consistent their party communication is.

The results indicate that our exploratory approach is useful to study how political parties profile themselves on Twitter and which issue strategies are at play. Second, our method allows to analyze communication of individual politicians, contributing to classical literature on party unity and party discipline. A comparison of our three methods shows a clear trade-off between interpretability and discriminative power. The expert-driven approach is insightful at the general issue level, but requires accurate upfront issue classification, which is hard to achieve. Recent advances in data-enhanced dictionaries, deep learning, transfer learning and semi-supervised learning offer exciting avenues for political text classification, but are increasingly complex and computationally demanding. Additionally, we find that a lot of information is lost by trying to categorize political communication on Twitter into predefined policy issues. In contrast, the data-driven approaches do not require upfront text classification and offer much more fine-grained insights at the event and even stylistic level of communication, but that comes at the expense of

interpretability at the issue level. A combination of all three methods simultaneously provides the best insights. This study demonstrates the benefits of an exploratory research approach for political communication research on social media.

In **Chapter 4** we study political communication from a network perspective. A lack of comparative empirical research in this domain leaves social scientists with little knowledge on the role that contextual factors play in the formation of social relations and network polarization. We collect one year of Twitter data from all members of parliament and government in 12 countries. Social network analysis is applied to analyze the relation between network properties of the parliamentary Twitter networks and the political system and democratic functioning of the countries. Secondly, we analyze the inter-party communication and its link to party ideology.

According to the results, consensual democracies are characterized by more dense parliamentary relations compared to majoritarian systems, but also by higher hierarchy and fragmentation. Furthermore, parliaments with a high effective number of parties are more cooperative, resulting in higher inter-party relations. The retweets network is most polarized or fragmented, while politicians engage more often in inter-party interactions in the followers and mentions network. The empirical evidence of the relationship between institutional context and network behavior can provide an explanation for contradictory research findings on network behavior and polarization suggested in previous (single-country) studies. This highlights the importance of including institutional context in social media research and the need for more comparative research to truly understand (online) behavior and the influence of social media in politics.

A second vital change to the political information flow is that social media enables citizens to self-select and distribute information. This feature fuels concerns that social media may contribute to growing political polarization. In **Part II**, we study the political behavior of the electorate on Facebook, as Facebook is by far the most popular social media platform among the general global population. Growing (perceived) political polarization invites speculation about the extent to which political polarization affects every aspect of our daily life. However, detailed information on individuals' lifestyle preferences is very difficult to collect, which complicates empirical and comparative studies in this domain. In **Chapter 5** we explore the potential of Facebook Likes to complement traditional survey data and study the interrelation between ideology and lifestyle choices. We collect a unique set of Facebook Likes and survey data from more than 6,500 participants in Belgium, and infer the political and ideological preference of our respondents. Additionally, we analyze the relatedness of different Facebook categories (e.g. movies, music, food, etc.) to ideology.

The results indicate that non-political Facebook Likes are indicative of political preference and are useful to describe voters in terms of common interests, cultural preferences, and lifestyle features. Moreover, some aspects of our social lives are

more connected to political preference than others. The reasons why this is the case are interesting follow-up questions for political and social scientists. In contrast with most previous research on lifestyle politics in the (polarized) two-party system in U.S., we examined lifestyle politics in a multi-party system. Facebook Likes appear to be less predictive for center voters and for traditional political parties, compared to voters with a more outspoken ideological position. Facebook Likes offer a rich source of information about individuals' revealed social and lifestyle preferences, at a resolution that would be difficult to attain with traditional survey techniques. Combining observational data with survey data helps us to understand the relationship between lifestyle preferences and politics at a fine-grained level. To reach insights across countries and time periods, future research needs to be comparative or at least similar in terms of data and methods

Therefore, a similar study is performed in the United States. In **Chapter 6** we utilize Facebook Like data to test whether polarization permeates society or if it is limited to strictly political domains and politically active individuals. We combine survey and Facebook Like data from more than 1,200 respondents in the United States in 2016. We analyze the network structure of political, news and lifestyle pages and calculate the average ideology and homogeneity of the audience per Facebook page. We observe that polarization is present in page categories that are somewhat related to politics — such as opinion leaders, political news sources, and topics related to identity and religion — but it does not appear to have strong influence in other domains, including sports, food, and music. On the individual level, we find that people with higher political news interest, stronger ideological predispositions, and an age of over 65 are more likely to endorse ideologically homogeneous pages across categories.

Finally, **Chapter 7** compares polarization in the two-party American system to the multi-party Belgian case. We find ideological divides in political domains to be much less outspoken in Belgium, and overall, more individuals seek out heterogeneous political spaces. Again, we find no evidence of high polarization in lifestyle domains. This finding seems to be in contrast to the earlier conclusion that lifestyle preferences are predictive for ideology and political leaning. For the majority of pages, we find that polarization is low, and a single Facebook page would provide very little evidence of a user's ideological beliefs. However, due to interaction effects, a combination of lifestyle preferences is very powerful in distinguishing left from right voters.

Our analysis based on Facebook Likes indicate a relationship between lifestyle and ideological identity but nuance the narrative of widespread polarization across lifestyle sectors. Nonetheless, these studies are snapshots in time and cannot rule out the possibility that lifestyle polarization is gradually increasing over time, as it is often claimed. Moreover, we believe it is likely that through its social signaling function, Facebook use is amplifying associations between lifestyle preferences and political identities. For instance, we uncovered a larger ideological divide in the news



diets on Facebook compared to website visits, which we argue is to be explained by this social signaling function. In any case, we recognize that mitigating polarization in both the online and offline world is an important challenge for society. In that respect, our results have important implications by providing insights into which individuals are more likely to seek out homogeneous political spaces and which domains offer most cross-cutting interactions.

## 8.2 The future of computational social media research

The amount of available data generated by our online footprints is likely to continue growing, as will the algorithmic possibilities to process these massive amounts of data. These two trends offer tremendous opportunities for the study of human behavior at unprecedented scale and detail. The different studies in this thesis confirm the potential of social media data for social and political science. Yet, as stated at the beginning of this thesis, it is not self-evident to convert this enormous empirical potential into valuable insights about human behavior. Big data can also lead to big mistakes if we do not follow a rigorous research approach: “To err is human, but to really foul things up you need a computer.”<sup>1</sup> So how do we continue to ask relevant questions about social behavior and the impact of social media on our society? Where do we find adequate online data and how do we apply rigorous methods to find answers to our questions? And finally, how do we do this in a transparent, responsible, and ethical manner?

**GOAL** The research goals of social and computer science differ. With machine learning we can predict age, gender, political affiliation, and personality from observational data, but without a clear research goal it is not necessarily relevant—or even ethically responsible—to do so. If we aim to understand social life, we need social science theory to pose meaningful questions in the first place. On the other hand, prediction tools are not always the most suitable method to answer a specific question. For example, when interested in the proportion of tweets that contain impolite language, should we base analysis on a representative subsample of the data that is manually coded and of high quality, or apply machine learning to classify the full sample? Furthermore, prediction methods are not suitable to find causal links, but they are very valuable for exploratory research by finding patterns in the data. A good understanding of these fundamental differences in explanation and prediction methods is needed. Interdisciplinary teams with social science and computer science expertise are the most qualified to answer meaningful questions with adequate research methods and reap the benefits of computational social science. As argued by Lazer et al. (2020), such interdisciplinary endeavors should be supported at universities, either by encouraging collaboration between

---

<sup>1</sup> The origin of this adage is unclear, some attribute it to biologist Paul Ehrlich, others to Alexander Pope, Senator Soaper, Bill Vaughan, or even Agatha Christie (Quote Investigator, 2021)

disciplines or by integrating elements of both disciplines into training social or computer scientists.

Second, as argued several times throughout this thesis, research on social media and politics is dominated by the United States, which is a two-party and highly polarized system. Often, these findings are not tested in the context of a multi-party system, with much more subtle ideological differences between parties. If we truly want to understand the mechanisms underlying political behavior and how this is influenced by social media we need to reach insights that travel across countries and time periods. To do so, future research needs to be comparative or at least more similar in terms of data and methods.

Finally, computational social scientists have the expertise to add to the debate about the role of social media in our society and can create an impact beyond a journal publication. Does social media influence our election results and democracy; does it increase political polarization, misinformation, online hate speech, and extremism; and how should we prevent this? Collaboration between scholars, policy makers, civil society, journalists, and political actors are needed to solve these stringent problems for humanity.

**DATA** Although the amount of observational data that is available in the digital world is unprecedented, researchers rely on the data sources to access this data. Twitter provides a free API that allows researchers to collect Twitter data. While the amount of data that could be collected with the free API was limited until shortly, Twitter has now opened up their full tweet archive to academic researchers in the new free Academic Research product track (Tornes and Trujillo, 2021). Facebook on the other hand, has restricted data collection for academic research through the APIs of Facebook, Instagram, and other platforms it owns after the Cambridge Analytica scandal. While this intervention certainly is positive for the privacy protection of their users, it is also locking out third parties and diminishing transparency of the platform. A promising avenue is the industry-academic partnership of Social Science One.<sup>2</sup> This partnership was designed to provide researchers with funding and data to study the effects of social media on democracy. Facebook released a large dataset of 38 million URLs to Social Science One to facilitate research on misinformation.<sup>3</sup> Because of the risk of re-identification, releasing data at this scale in a privacy-preserving way is challenging and took Facebook more than two years. Despite Facebook's efforts to collaborate with researchers through Social Science One, some scholars question whether this initiative will provide sufficient support for free and independent scientific research (Bruns et al., 2018).

---

<sup>2</sup> <https://socialscience.one/>

<sup>3</sup> <https://socialscience.one/blog/unprecedented-facebook-urls-dataset-now-available-research-through-social-science-one>

In his book *The Hype Machine*, Sinan Aral mentions three essential steps to enable scientific research using social media data. First, we need technical solutions to the transparency paradox. We are asking platforms to protect user privacy and be transparent with their methods and data at the same time. To release data anonymously, noise can be added to it. However, the more noise is added, the less informative the data becomes. Designing differential privacy that maximizes the utility of the data for research while protecting privacy is one of the next big challenges in computer science. Second, policy makers and social platforms should create safe harbors in which researchers can access and analyze sensitive data. The amount and type of data they can access is limited, as well as the analysis they can perform. The creators of Social Science One encourage the European Commission and the Federal Trade Commission to facilitate such safe harbors.<sup>4</sup> Third, platforms must make a commitment to make accurate and representative data available for scientific research. Twitter has taken this responsibility to heart with the new free Academic Research product track. Facebook provides researchers access to the political ads on their platform through the Ads Library<sup>5</sup> and has released the URL dataset to Social Science One. We have focused on Twitter and Facebook throughout this thesis but of course other platforms are equally –if not more– relevant, such as Youtube and TikTok, and should commit to make data available for research as well. TikTok is becoming more and more popular –especially among the younger population– and politically relevant (Medina Serrano et al., 2020). Research on this relatively new platform should not lag behind.

The challenges just discussed also have consequences for reproducibility of academic research. If sensitive data can only be accessed under strict conditions, it cannot be shared with other researchers to replicate results. The scientific community should rethink the standards of research transparency. As a solution, making aggregated results and code publicly available is often sufficient in the case of sensitive data.

Due to restrictions on sharing data, little coordinated efforts exist to collect and centralize publicly available data. For example, it is allowed to share the IDs of tweets collected through the Twitter API and large Twitter archives do exist,<sup>6</sup> yet most researchers keep collecting similar data in parallel and do not share their collections (or share it on different platforms such as Github). On the positive side, several initiatives have come to live over the past years that help researchers without technical know-how to collect and/or analyze social media data. For example, the Social Media Analysis Toolkit (SMAT)<sup>7</sup> helps to facilitate the analysis and visualization of larger trends on a variety of platforms.

<sup>4</sup> <https://socialscience.one/blog/public-statement-european-advisory-committee-social-science-one>

<sup>5</sup> <https://www.facebook.com/ads/library/>

<sup>6</sup> such as DocNow (<https://catalog.docnow.io/>) and TweetSets (<https://tweetsets.library.gwu.edu/>)

<sup>7</sup> <https://www.smat-app.com/about>

**METHODS** Machine learning models are becoming increasingly accurate, and ever more complex. The increasing complexity of prediction algorithms is accompanied by an equal effort in trying to explain the decisions made by these black boxes. Humans need to understand the model in order to trust its decisions and apply and improve it in practice. Interpretability of the model helps to detect and correct bias and ensures us that meaningful variables infer the output (Arrieta et al., 2020). Within social science, where transparency and interpretability of the results are key, we cannot simply apply black box models.

As Rudin (2019) advocates we should apply inherently interpretable models in the first place instead of trying to explain black box models. However, in some situations the increase in accuracy might justify the increase in complexity. Moreover, when learning from behavioral and textual data, the complexity of the model arises from the nature of the data itself, not from the learning technique. The high-dimensionality and sparsity of the data combined with many relevant features are making even linear models hard to interpret, as we need to examine thousands of coefficients (Ramon et al., 2021). For these reasons, the recent advances in explainable machine learning will offer fruitful opportunities for computational social science to extract relevant insights from complex data and methods. Again, this requires interdisciplinary endeavors as these model explanations need to be understandable and relevant for social scientists in the first place.

Although we have focused on behavioral and textual data, I should note that audio-visual data is on the rise on social media. A picture is worth a thousand words: images trigger stronger emotional reactions than text and are more mobilizing (Casas and Williams, 2019). Therefore, the use of images and video is popular on social media and worth investigating. Processing such information demands advanced speech and image recognition techniques. These methods are very complex, once again pointing towards the relevance of explainability in machine learning.

**ETHICAL AND LEGAL CONSIDERATIONS** We have touched upon several ethical issues related to online collection, storage, and use of human subjects' data. Regulation (GDPR) provides us with minimum requirements, but ethical responsibilities go beyond what is legally obliged. Yet, the variety in data sources, research topics, and methodological approaches complicates the draft of universally applicable ethical guidelines. Rather, all researchers engaging with (data from) human subjects, have the ethical responsibility to minimize potential harm. Transparent communication about ethical research design in scientific journals could benefit the ethical debate and framework for online human subject research. I encourage all researchers and journals to include such discussion in their work. Research ethics should be an ongoing debate in the scientific field and an integral part of each scholars' training.

### 8.3 The future of social media

How far are we from the Black Mirror episode “Nosedive”, where we’re completely dependent upon social media and our life is determined by likes, followers and ratings? It is clear that social media has drastically changed our lives (for better or worse) and that research is needed to understand its impact on our society and democracy. I want to end this thesis with a brief outlook on the future of social media and our role in achieving a positive coexistence with this technology. I can only applaud initiatives like the Netflix documentary “The Social Dillema” that bring the perils of social media to the attention of the general public. I think awareness and education are key towards a better future. At the same time, “The Social Dillema” misses scientific substantiation and puts the blame very one-sidedly on the technology and its creators. While they certainly have a responsibility in both the cause and the solution to the problem, it is an interplay of several stakeholders that has got us into the current situation. And it is the same mix of elements that can get us out again: platforms, public sector, technological innovation, and personal responsibility (Hilbert, 2020).<sup>8</sup>

To illustrate how we need all four strategies simultaneously to improve our situation, let’s consider an excellent example described by Aral (2020), regarding the antitrust case against Facebook. Facebook’s monopoly position can be considered as a market failure, and some reason that breaking up big tech companies will force them to compete in protecting our social values, like securing privacy, and reducing misinformation and harmful speech. Antitrust law can break up monopolies if they harm consumers in an uncompetitive market. However, if we break up Facebook,<sup>9</sup> its place will soon be taken by a successor. The social media market itself is subject to strong network effects: the platform’s value increases when more users are on the platform. If more of your friends are on Facebook, you have a stronger incentive to also use Facebook. A solution to this problem is interoperability and data portability. Similar to the telecommunications industry where you can call from one network provider to the other and transfer your number when you want to change to another provider, imagine you can just move your social graph to a different platform. Regulators could demand such interoperability to ensure competition. This would require collaboration among platforms to design interoperable protocols and massive technological innovation, as interoperability and data portability in social networks is much more complex than in the telecom industry. Eventually, if we were to achieve a truly competitive market, the consumers still need to define the values they want platforms to deliver to them. If consumers do not value privacy protection over exciting conspiracy theories, a competitive market won’t bring us one step closer to a solution for data abuse, fake news, hate speech, and polarization. Let’s now have a closer look at each of these stakeholders’ responsibilities and power.

<sup>8</sup> Aral (2020) calls it money, law, code and norms but essentially refers to similar concepts.

<sup>9</sup> I mean the social platform Facebook, not the company Facebook which also consists of Whatsapp and Instagram. Uncoupling Whatsapp and Instagram is a slightly different story.

**PLATFORMS** In the social media industry, attention is money. Platforms and their algorithms are designed to keep us engaged and sell our attention to advertisers (Myllylahti, 2018). The social media logic, or the processes through which these platforms channel social traffic, is tuned to show us that information or content that engages us most. As consumers are attracted to novel and outrageous news (Vosoughi et al., 2018) and content that is congruent with their own beliefs (Sikder et al., 2020) it is not hard to imagine how the social media logic fuels the spread of false news and polarization. But social platforms can also adapt their logic to counteract the spread of misinformation and recommend cross-cutting and diverse information. Although, on the other side, some researcher argue that showing opposing content may aggravate sectarianism rather than reducing it (Finkel et al., 2020). In Part II, we discovered that some lifestyle categories are more likely to unite people with opposing political views. Facebook pages with a diverse audience could be recommended more often to dampen political polarization and maintain social harmony. Just recently, Facebook announced it will stop recommending political groups to users (BBC, 2021). Similarly, friend suggestions are mostly based on common friends and similar interests, and thus have the tendency to aggravate political homophily. Next to content recommendations, social networks could reevaluate their friend recommendations to encourage more diverse friend networks.

Furthermore, platforms could also actively invest in labeling fake news or nudging people into thinking about the accuracy of the information they read (Aral, 2020). In several survey experiments and a field experiment on Twitter, Pennycook et al. (2020) found that priming participants to think about accuracy does indeed increase the quality of news that people subsequently share. To avoid harmful speech regulators and society have pressured platforms to take responsibility and moderate their content. Which then in turn raises questions about how far platforms can go in deciding what is allowed and what is not, and whether this should not be more regulated. Banning the account of Donald Trump was one bridge too far according to some and just about time according to others (Breton, 2021).

To conclude, I think most social platforms are starting to take their responsibility—under large societal pressure—in avoiding the perils they created. During the 2020 U.S. presidential election, many platforms took action to combat misinformation and avoid interference. For example, by banning political ads or labeling or removing false election claims.<sup>10</sup> After the Cambridge Analytica incident and pressured by GDPR regulation, Facebook and others have thoroughly revised their privacy policy. On the other hand, it still remains their business model to know you as well as possible.

**PUBLIC SECTOR** To mitigate market failures, and protect our privacy and free speech, laws and regulation can impose restrictions on the platforms' economically

<sup>10</sup> See <https://www.eipartnership.net/policy-analysis/platform-policies>

optimized design choices. GDPR regulation in Europe and State law in the U.S. protect our personal data that is used by social media platforms. Importantly, as we regulate privacy we should not preclude the possibilities for independent research. In Europe, the Digital Services Act (DSA) and the Digital Markets Act (DMA) have the goals to “create a safer digital space in which the fundamental rights of all users of digital services are protected” and “to establish a level playing field to foster innovation, growth, and competitiveness, both in the European Single Market and globally” (European Commission, 2021). The regulation will oblige platforms to disclose how their algorithms work, and how content is moderated. Similarly, in the U.S., the Biden administration is preparing reforms to antitrust law meant to rein in the biggest tech companies. During his campaign, Biden also announced he would reform Section 230 of the Communications Decency Act, to set out rules for content moderation (Reardon, 2021).

Next to regulation, an essential task of the public sector is to create awareness and organize education to inform citizens about the possible dangers of their behavior on social media. Training in critical thinking and media literacy can reduce the spread of misinformation. Research has found that fake news dissemination in the U.S. is especially prevalent among users over the age of 65 (Guess et al., 2019b). Social media education should not only be part of the curriculum in schools but we should also try to reach the older generation.

**TECHNOLOGICAL INNOVATION** Just as technology is at the root of many of today’s problems with social media, they can play an important role in creating the solutions. Based on the content and network-based features of news articles or messages, state-of-the-art data mining algorithms are capable of automatically detecting false news (Shu et al., 2017) and harmful speech (Schmidt and Wiegand, 2017). This can greatly assist the platforms in labeling or removing these types of messages, as the content created on most platforms has become too extensive to manually oversee. Smart recommendation algorithms can create exposure to a wide range of views. A team of researchers developed a new algorithm that increases the diversity of exposure on social networks, while still ensuring that content is widely shared (Aslay et al., 2018). Lastly, advancements in privacy-preserving techniques could help solve the transparency paradox by protecting private information while disclosing essential information for independent research.

**PERSONAL RESPONSIBILITY** Fake news travels faster than the truth because humans, not robots, are more likely to spread it (Vosoughi et al., 2018). Echo chambers arise mainly because of individuals’ choices rather than algorithmic ranking (Bakshy et al., 2015). In fact, they might even be greater in offline social networks, where exposure to diverse views is often more rare (Guess et al., 2018).<sup>11</sup> Again, education

<sup>11</sup> Of course these are just a few examples of studies that nuance the effects of social media, numerous counterexamples exist as well. The point is that we should not minimize the impact of our own behavior.



## CONCLUSIONS

can help us to make more conscious online decisions. Never before have we had such opportunities to actively seek out cross-cutting information, to connect with a broad network, and to engage in discussion. Recognizing that polarization is a larger societal trend, that is reinforced and catalyzed, but not caused by social media use (Moeller et al., 2018), will help us design effective strategies to combat it. Finally, we control social media just as much as it controls us. We are the customers and the products that drive most of its design and so we need to define the values we want platforms to deliver to us.

In conclusion, to build a better online future we need a complex mix of strategies. As humans, we should not underestimate the power we have in controlling the technology we have created. Nonetheless, it will be a continuous, uncertain, and complicated endeavor, and we need research to guide us along the way. This will bring us one step closer to the dream of a truly connected world.



## Bibliography

---

- Abramowitz, A. I. and Saunders, K. L. (2008). Is polarization a myth? *The Journal of Politics*, 70(2):542–555.
- Adamic, L. A. and Glance, N. (2005). The political blogosphere and the 2004 us election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43.
- Ahler, D. J. and Sood, G. (2018). The parties in our heads: Misperceptions about party composition and their consequences. *The Journal of Politics*, 80(3):964–981.
- Al-Ayyoub, M., Rabab’ah, A., Jararweh, Y., Al-Kabi, M. N., and Gupta, B. B. (2018). Studying the controversy in online crowds’ interactions. *Applied Soft Computing*, 66:557–563.
- Alford, J. R., Hatemi, P. K., Hibbing, J. R., Martin, N. G., and Eaves, L. J. (2011). The politics of mate choice. *The Journal of Politics*, 73(2):362–379.
- Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–36.
- Allcott, H., Gentzkow, M., and Yu, C. (2019). Trends in the diffusion of misinformation on social media. *Research & Politics*, 6(2):2053168019848554.
- Amarasingam, A. and Argentino, M. (2020). The qanon conspiracy theory: A security threat in the making? *CTC Sentinel*, 13:37–44.
- Andeweg, R. B. and Thomassen, J. (2011). Pathways to party unity: Sanctions, loyalty, homogeneity and division of labour in the dutch parliament. *Party Politics*, 17(5):655–672.
- Aral, S. (2020). *The Hype Machine: How Social Media Disrupts Our Elections, Our Economy, and Our Health-and how We Must Adapt*. Currency.
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., et al. (2020). Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58:82–115.
- Aslay, C., Matakos, A., Galbrun, E., and Gionis, A. (2018). Maximizing the diversity of exposure in a social network. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 863–868. IEEE.
- Bach, R. L., Kern, C., Amaya, A., Keusch, F., Kreuter, F., Hecht, J., and Heinemann, J. (2019). Predicting voting behavior using digital trace data. *Social Science Computer Review*, 37(5):591–683.
- Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.

- Bakos, Y., Marotta-Wurgler, F., and Trossen, D. R. (2014). Does anyone read the fine print? consumer attention to standard-form contracts. *The Journal of Legal Studies*, 43(1):1–35.
- Bakshy, E., Messing, S., and Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239):1130–1132.
- Baldassarri, D. and Gelman, A. (2008). Partisans without constraint: Political polarization and trends in american public opinion. *American Journal of Sociology*, 114(2):408–446.
- Barbaro, M., Zeller, T., and Hansell, S. (2006). A face is exposed for aol searcher no. 4417749. *New York Times*, 9(2008):8.
- Barberá, P. (2015). Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political analysis*, 23(1):76–91.
- Barberá, P., Boydston, A. E., Linn, S., McMahon, R., and Nagler, J. (2019). Automated text classification of news articles: A practical guide. *Political Analysis*, pages 1–24.
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., and Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, 26(10):1531–1542. PMID: 26297377.
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., and Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science*, 26(10):1531–1542.
- Barberá, P. and Steinert-Threlkeld, Z. C. (2020). How to use social media data for political science research. *The SAGE Handbook of Research Methods in Political Science and International Relations*, page 404.
- Barclay, F. P., Pichandy, C., Venkat, A., and Sudhakaran, S. (2015). India 2014: Facebook ‘Like’ as a Predictor of Election Outcomes. *Asian Journal of Political Science*, 23(2):134–160.
- Barocas, S. and Nissenbaum, H. (2014). Big data’s end run around anonymity and consent. *Privacy, big data, and the public good: Frameworks for engagement*, 1:44–75.
- BBC (2021). Facebook to stop recommending civic and political groups.
- Belford, G. and Tucker, A. (2020). Computer science.
- Benkler, Y., Faris, R., and Roberts, H. (2018). *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford University Press.
- Bennett, W. L. and Segerberg, A. (2012). The logic of connective action: Digital media and the personalization of contentious politics. *Information, communication & society*, 15(5):739–768.
- Berelson, B. R., Lazarsfeld, P. F., McPhee, W. N., and McPhee, W. N. (1954). *Voting: A study of opinion formation in a presidential campaign*. University of Chicago Press.
- Bértoa, F. C. (2020). Database on who governs in europe and beyond, psgo. university of nottingham.
- Bimber, B. and Davis, R. (2003). *Campaigning online: The Internet in US elections*. Oxford University Press.
- Bishop, B. (2009). *The big sort: Why the clustering of like-minded America is tearing us apart*. Houghton Mifflin Harcourt.
- Blei, D. M. and Lafferty, J. D. (2006). Dynamic topic models. In *Proceedings of the 23rd international conference on Machine learning*, pages 113–120. ACM.

- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008.
- Boerboom, C. (2020). Cambridge analytica: The scandal on data privacy.
- Bond, R. and Messing, S. (2015). Quantifying social media’s political space: Estimating ideology from publicly revealed preferences on facebook. *American Political Science Review*, 109(1):62–78.
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., and Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298.
- Borge Bravo, R. and Esteve Del Valle, M. (2017). Opinion leadership in parliamentary twitter networks: A matter of layers of interaction? *Journal of Information Technology & Politics*, 14(3):263–276.
- Boutyline, A. and Willer, R. (2017). The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, 38(3):551–569.
- Bovet, A. and Makse, H. A. (2019). Influence of fake news in twitter during the 2016 us presidential election. *Nature communications*, 10(1):1–14.
- Boxell, L., Gentzkow, M., and Shapiro, J. M. (2017). Greater internet use is not associated with faster growth in political polarization among us demographic groups. *Proceedings of the National Academy of Sciences*, 114(40):10612–10617.
- Boyd, D. M. and Ellison, N. B. (2007). Social network sites: Definition, history, and scholarship. *Journal of computer-mediated communication*, 13(1):210–230.
- Breton, T. (2021). Capitol hill — the 9/11 moment of social media.
- Bruns, A., Bechmann, A., Burgess, J., Chadwick, A., Clark, L. S., Dutton, W. H., Ess, C. M., Gruz, A., Halford, S., Hofmann, J., et al. (2018). Facebook shuts the gate after the horse has bolted, and hurts real research in the process. *Internet Policy Review*.
- Buntain, C., McGrath, E., Golbeck, J., and LaFree, G. (2016). Comparing social media and traditional surveys around the boston marathon bombing. In *Proceedings of the 6th Workshop on ‘Making Sense of Microposts’*, pages 34–41. WWW2016.
- Burkart, N. and Huber, M. F. (2020). A survey on the explainability of supervised machine learning. *arXiv preprint arXiv:2011.07876*.
- Campbell, A., Converse, P. E., Miller, W. E., and Stokes, D. E. (1960). *The American Voter*. University of Chicago Press.
- Casas, A. and Williams, N. W. (2019). Images that matter: Online protests and the mobilizing role of pictures. *Political Research Quarterly*, 72(2):360–375.
- Ceron, A., Curini, L., and Iacus, S. (2015). Using social media to fore-cast electoral results: A review of state-of-the-art. *Italian Journal of Applied Statistics*, 25:237–259.
- Chaisty, P. (2008). The legislative effects of presidential partisan powers in post-communist russia. *Government and Opposition*, 43(3):424–453.
- Chang, J., Boyd-Graber, J., Wang, C., Gerrish, S., and Blei, D. M. (2009). Reading tea leaves: How humans interpret topic models. In *Neural information processing systems*, volume 22, pages 288–296. Citeseer.

- Cherepnalkoski, D. and Mozetič, I. (2016). Retweet networks of the european parliament: evaluation of the community structure. *Applied network science*, 1(1):1–20.
- Chiu, S.-I. and Hsu, K.-W. (2018). Predicting political tendency of posts on facebook. In *Proceedings of the 2018 7th International Conference on Software and Computer Applications*, pages 110–114. ACM.
- Christensen, R. K., Rakhimkulov, E. R., and Wise, C. R. (2005). The ukrainian orange revolution brought more than a new president: What kind of democracy will the institutional changes bring? *Communist and Post-Communist Studies*, 38(2):207–230.
- Clark, J. and Provost, F. (2015). Dimensionality reduction via matrix factorization for predictive modeling from large, sparse behavioral data.
- Clauset, A., Newman, M. E., and Moore, C. (2004). Finding community structure in very large networks. *Physical review E*, 70(6):066111.
- Clement, J. (2020a). Distribution of facebook users worldwide as of october 2020, by age and gender. *Statistica*.
- Clement, J. (2020b). Global social networks ranked by number of users 2020. *Statistica*.
- Clement, J. (2020c). Number of global social network users 2017-2025. *Statistica*.
- CNN (2021a). In pictures: Trump’s second impeachment.
- CNN (2021b). Trump is losing his powerful social media machine. now what?
- CNN (2021c). Us capitol secured, 4 dead after rioters stormed the halls of congress to block biden’s win.
- Cohen, J. (2013). *Statistical power analysis for the behavioral sciences*. Academic press.
- Colnerić, N. and Demsar, J. (2018). Emotion recognition on twitter: Comparative study and training a unison model. *IEEE transactions on affective computing*.
- Conover, M., Ratkiewicz, J., Francisco, M. R., Gonçalves, B., Menczer, F., and Flammini, A. (2011). Political polarization on twitter. *ICWSM*, 133:89–96.
- Conover, M. D., Gonçalves, B., Flammini, A., and Menczer, F. (2012). Partisan asymmetries in online political activity. *EPJ Data Science*, 1(1):6.
- Conover, P. J. and Feldman, S. (1981). The Origins and Meaning of Liberal/Conservative Self-Identifications. page 30.
- Contreras-Piña, C. and Ríos, S. A. (2016). An empirical comparison of latent semantic models for applications in industry. *Neurocomputing*, 179:176–185.
- Cook, J. M. (2016). Are american politicians as partisan online as they are offline? twitter networks in the us senate and maine state legislature. *Policy & Internet*, 8(1):55–71.
- Coppersmith, G., Dredze, M., and Harman, C. (2014). Quantifying mental health signals in twitter. In *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*, pages 51–60.
- Council of Europe (1950). European convention for the protection of human rights and fundamental freedoms.
- Cranmer, S. J. and Desmarais, B. A. (2017). What can we learn from predictive modeling? *Political Analysis*, 25(2):145–166.
- Croarkin, C., Tobias, P., Zey, C., et al. (2002). *Engineering statistics handbook*. NIST iTL.
- Curtin, F. and Schulz, P. (1998). Multiple correlations and bonferroni’s correction. *Biological psychiatry*, 44(8):775–777.

- Dalton, R. J. (2016). The potential of big data for the cross-national study of political behavior. *International Journal of Sociology*, 46(1):8–20.
- Dalton, R. J., McAllister, I., and Wattenberg, M. P. (2000). The consequences of partisan dealignment. In *Parties without partisans: Political change in advanced industrial democracies*, pages 37–63. Oxford University Press Oxford.
- Dalton, R. J. and Wattenberg, M. P. (2002). *Parties without partisans: Political change in advanced industrial democracies*. Oxford University Press.
- Damore, D. F. (2005). Issue convergence in presidential campaigns. *Political Behavior*, 27(1):71–97.
- Daniels, J. (2018). The algorithmic rise of the “alt-right”. *Contexts*, 17(1):60–65.
- David, E., Zhitomirsky-Geffet, M., Koppel, M., and Uzan, H. (2016). Utilizing Facebook pages of the political parties to automatically predict the political orientation of Facebook users. *Online Information Review*, 40(5):610–623.
- De Cnudde, S. (2017). *An Analysis of Wide and Deep Classification Techniques for Human Behavioral Data*.
- De Cnudde, S., Martens, D., Evgeniou, T., and Provost, F. (2017). A benchmarking study of classification techniques for behavioral data.
- De Cnudde, S., Moeyersoms, J., Stankova, M., Tobback, E., Javalý, V., and Martens, D. (2019). What does your facebook profile reveal about your creditworthiness? using alternative data for microfinance. *Journal of the Operational Research Society*, 70(3):353–363.
- De Sio, L. and Lachat, R. (2020). Making sense of party strategy innovation: challenge to ideology and conflict-mobilisation as dimensions of party competition. *West European Politics*, 43(3):688–719.
- DellaPosta, D. (2020). Pluralistic collapse: The “oil spill” model of mass opinion polarization. *American Sociological Review*, page 0003122420922989.
- DellaPosta, D., Shi, Y., and Macy, M. (2015). Why do liberals drink lattes? *American Journal of Sociology*, 120(5):1473–1511.
- Delwit, P., Deschouwer, K., Hooghe, M., Baudewijns, P., and Walgrave, S. (2015). *De kiezer ontcijferd. Over stemgedrag en stemmotivaties*. LannooCampus.
- Depauw, S. and Martin, S. (2009). Legislative party discipline and cohesion in comparative perspective. *Intra-party politics and coalition governments*, 103120:103–120.
- Desmet, K., Ortuño-Ortín, I., and Wacziarg, R. (2017). Culture, ethnicity, and diversity. *American Economic Review*, 107(9):2479–2513.
- Dollarhide, M. (2020). Social media definition. *Investopedia*.
- Druckman, J. N., Green, D. P., Kuklinski, J. H., and Lupia, A. (2006). The growth and development of experimental research in political science. *American Political Science Review*, 100(4):627–635.
- Druckman, J. N., Peterson, E., and Slothuus, R. (2013). How elite partisan polarization affects public opinion formation. *American Political Science Review*, pages 57–79.
- Dun, L., Soroka, S., and Wlezien, C. (2020). Dictionaries, supervised learning, and media coverage of public policy. *Political Communication*, pages 1–19.
- Dvir-Gvirsman, S. (2017). Media audience homophily: Partisan websites, audience identity and polarization processes. *New media & society*, 19(7):1072–1091.

- Eady, G., Nagler, J., Guess, A., Zilinsky, J., and Tucker, J. A. (2019). How many people live in political bubbles on social media? evidence from linked survey and twitter data. *Sage Open*, 9(1):1–22.
- EIU (2020). Democracy index 2019. a year of democratic setbacks and popular protest.
- Enikolopov, R., Makarin, A., and Petrova, M. (2020). Social media and protest participation: Evidence from russia. *Econometrica*, 88(4):1479–1514.
- Eranti, V. and Lonkila, M. (2015). The social significance of the facebook like button. *First Monday*, 20.
- Esteve Del Valle, M. and Borge Bravo, R. (2018). Echo chambers in parliamentary twitter networks: The catalan case. *International journal of communication*, 12:21.
- European Commission (2021). The digital services act package.
- European Parliament and Council (2016). Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation). *Official Journal of the European Union*.
- Evans, J. H., Bryson, B., and DiMaggio, P. (2001). Opinion polarization: Important contributions, necessary limitations. *American Journal of Sociology*, 106(4):944–959.
- Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., Mason, L., McGrath, M. C., Nyhan, B., Rand, D. G., et al. (2020). Political sectarianism in america. *Science*, 370(6516):533–536.
- Fiorina, M. P., Abrams, S. A., and Pope, J. C. (2008). Polarization in the american public: Misconceptions and misreadings. *The Journal of Politics*, 70(2):556–560.
- Fischer, M. (2011). Social network analysis and qualitative comparative analysis: Their mutual benefit for the explanation of policy network structures. *Methodological Innovations Online*, 6(2):27–51.
- Flach, P. A., Hernández-Orallo, J., and Ramirez, C. F. (2011). A coherent interpretation of auc as a measure of aggregated classification performance. In *ICML*.
- Francis, J. G. and Payne, C. (1977). The use of the logistic model in political science: British elections, 1964-1970. *Political Methodology*, 4(3):233–270.
- Franzke, A., Bechmann, A., Zimmer, M., Ess, C., and the association of internet researchers (2020). Internet research: ethical guidelines 3.0.
- Freelon, D. (2018). Computational research in the post-api age. *Political Communication*, 35(4):665–668.
- Freeman, L. C. (1978). Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239.
- Fruchterman, T. M. and Reingold, E. M. (1991). Graph drawing by force-directed placement. *Software: Practice and experience*, 21(11):1129–1164.
- Funk, C. L. (1999). Bringing the candidate into models of candidate evaluation. *The Journal of Politics*, 61(3):700–720.
- Furnham, A. (1986). Response bias, social desirability and dissimulation. *Personality and Individual Differences*, 7(3):385–400.
- Gainous, J. and Wagner, K. M. (2013). *Tweeting to power: The social media revolution in American politics*. Oxford University Press.

- Garcia, D., Abisheva, A., Schweighofer, S., Serdült, U., and Schweitzer, F. (2015). Ideological and temporal components of network polarization in online political participatory media. *Policy & internet*, 7(1):46–79.
- Gates, S. and Podder, S. (2015). Social media, recruitment, allegiance and the islamic state. *Perspectives on Terrorism*, 9(4):107–116.
- Geng, L. and Hamilton, H. J. (2006). Interestingness measures for data mining: A survey. *ACM Computing Surveys (CSUR)*, 38(3):9.
- Gentzkow, M., Shapiro, J., Taddy, M., et al. (2016). Measuring polarization in high-dimensional data: Method and application to congressional speech.
- Giddens, A. (2013). *The consequences of modernity*. John Wiley & Sons.
- Giglietto, F. (2012). If Likes Were Votes: An Empirical Study on the 2011 Italian Administrative Elections. *SSRN Electronic Journal*.
- Goffman, E. (1967). *Interaction ritual: Essays on face-to-face interaction*.
- Goga, O., Venkatadri, G., and Gummadi, K. P. (2015). The doppelgänger bot attack: Exploring identity impersonation in online social networks. In *Proceedings of the 2015 internet measurement conference*, pages 141–153.
- Golbeck, J., Robles, C., Edmondson, M., and Turner, K. (2011). Predicting personality from twitter. In *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, pages 149–156. IEEE.
- González, R. J. (2017). Hacking the citizenry?: Personality profiling, ‘big data’ and the election of donald trump. *Anthropology Today*, 33(3):9–12.
- Gorton, W. A. (2016). Manipulating citizens: how political campaigns’ use of behavioral social science harms democracy. *New Political Science*, 38(1):61–80.
- Graham, T., Jackson, D., and Broersma, M. (2016). New platform, old habits? candidates’ use of twitter during the 2010 british and dutch general election campaigns. *New media & society*, 18(5):765–783.
- Grant, W. J., Moon, B., and Busby Grant, J. (2010). Digital dialogue? australian politicians’ use of the social network tool twitter. *Australian Journal of Political Science*, 45(4):579–604.
- Green, D. P., Palmquist, B., and Schickler, E. (2004). *Partisan hearts and minds: Political parties and the social identities of voters*. Yale University Press.
- Green-Pedersen, C. (2007). The growing importance of issue competition: The changing nature of party competition in western europe. *Political studies*, 55(3):607–628.
- Greenberg, J. and Jonas, E. (2003). Psychological motives and political orientation—the left, the right, and the rigid: comment on jost et al.(2003).
- Greene, T., Shmueli, G., Ray, S., and Fell, J. (2019). Adjusting to the gdpr: The impact on data scientists and behavioral researchers. *Big data*, 7(3):140–162.
- Grimmer, J. and Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political analysis*, 21(3):267–297.
- Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., and Lazer, D. (2019). Fake news on twitter during the 2016 us presidential election. *Science*, 363(6425):374–378.

- Gross, P. (2008). Dances with wolves: A meditation on the media and political system in the european union's romania. *Finding the right place on the map. Central and Eastern European media change in a global perspective*, pages 125–145.
- Gruzd, A. and Roy, J. (2014). Investigating political polarization on twitter: A canadian perspective. *Policy & internet*, 6(1):28–45.
- Guess, A., Munger, K., Nagler, J., and Tucker, J. (2019a). How accurate are survey responses on social media and politics? *Political Communication*, 36(2):241–258.
- Guess, A., Nagler, J., and Tucker, J. (2019b). Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science advances*, 5(1):eaau4586.
- Guess, A., Nyhan, B., Lyons, B., and Reifler, J. (2018). Avoiding the echo chamber about echo chambers. *Knight Foundation*.
- Guess, A. M. (2020). (almost) everything in moderation: New evidence on americans' online media diets.
- Guess, A. M., Nyhan, B., and Reifler, J. (2020). Exposure to untrustworthy websites in the 2016 us election. *Nature human behaviour*, 4(5):472–480.
- Gupta, V. and Hewett, R. (2020). Real-time tweet analytics using hybrid hashtags on twitter big data streams. *Information*, 11(7):341.
- Han, B. and Baldwin, T. (2011). Lexical normalisation of short text messages: Makn sens a #twitter. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1, HLT '11*, pages 368–378, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Hand, D. J. (2009). Measuring classifier performance: a coherent alternative to the area under the roc curve. *Machine learning*, 77(1):103–123.
- Hasan, M., Agu, E., and Rundensteiner, E. (2014). Using hashtags as labels for supervised learning of emotions in twitter messages. In *ACM SIGKDD workshop on health informatics*, New York, USA.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.
- Hetherington, M. and Weiler, J. (2018). *Prius or pickup?: How the answers to four simple questions explain America's great divide*. Houghton Mifflin.
- Hetherington, M. J. (2009). Putting polarization in perspective. *British Journal of Political Science*, pages 413–448.
- Higley, J. and Burton, M. G. (1989). The elite variable in democratic transitions and breakdowns. *American sociological review*, pages 17–32.
- Hilbert, M. (2020). Social media distancing: An opportunity to debug our relationship with our algorithms.
- Himmelboim, I., Smith, M. A., Rainie, L., Shneiderman, B., and Espina, C. (2017). Classifying twitter topic-networks using social network analysis. *Social Media+ Society*, 3(1):2056305117691545.
- Hollenbeck, C. R. and Kaikati, A. M. (2012). Consumers' use of brands to reflect their actual and ideal selves on facebook. *International Journal of Research in Marketing*, 29(4):395–405.
- Hooghe, M. and Dassonneville, R. (2018). A spiral of distrust: A panel study on the relation between political distrust and protest voting in belgium. *Government and Opposition*, 53(1):104–130.



- Hopkins, D. J. and King, G. (2010). A method of automated nonparametric content analysis for social science. *American Journal of Political Science*, 54(1):229–247.
- Howard, P. N., Duffy, A., Freelon, D., Hussain, M. M., Mari, W., and Maziad, M. (2011). Opening closed regimes: what was the role of social media during the arab spring? *Available at SSRN 2595096*.
- Huber, G. A. and Malhotra, N. (2017). Political homophily in social relationships: Evidence from online dating behavior. *The Journal of Politics*, 79(1):269–283.
- IDEA (2019). Electoral system design database.
- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., and Westwood, S. J. (2019). The origins and consequences of affective polarization in the united states. *Annual Review of Political Science*, 22:129–146.
- Iyengar, S., Sood, G., and Lelkes, Y. (2012). Affect, not ideology a social identity perspective on polarization. *Public opinion quarterly*, 76(3):405–431.
- Jackson, M. (2010). *Social and Economic Networks*.
- Jaiswal, V. and Agarwal, J. (2012). The Evolution of the Association Rules. *International Journal of Modeling and Optimization*.
- Jalali-Heravi, M. and Zaïane, O. R. (2010). A study on interestingness measures for associative classifiers. In *Proceedings of the 2010 ACM Symposium on Applied Computing*, pages 1039–1046. ACM.
- Johansson, E. (2019). Social media in political communication a substitute for conventional media ?
- Jones, E., Oliphant, T., Peterson, P., et al. (2001). SciPy: Open source scientific tools for Python.
- Jones, J. J., Bond, R. M., Bakshy, E., Eckles, D., and Fowler, J. H. (2017). Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 us presidential election. *PloS one*, 12(4):e0173851.
- Joseph, K. and Wihbey, J. (2019). Breaking news and younger twitter users: Comparing self-reported motivations to online behavior. In *Proceedings of the 10th International Conference on Social Media and Society*, pages 83–91.
- Jost, J. T., Federico, C. M., and Napier, J. L. (2009). Political ideology: Its structure, functions, and elective affinities. *Annual review of psychology*, 60:307–337.
- Jungherr, A. (2016). Twitter use in election campaigns: A systematic literature review. *Journal of information technology & politics*, 13(1):72–91.
- Jungherr, A. (2018). Normalizing digital trace data. *Digital discussions. How big data informs political communication*. Oxon: Routledge, pages 19–45.
- Jungherr, A. and Theocharis, Y. (2017). The empiricist’s challenge: Asking meaningful questions in political science in the age of big data.
- Junqué de Fortuny, E., Martens, D., and Provost, F. (2013). Predictive modeling with big data: is bigger really better? *Big Data*, 1(4):215–226.
- Jurafsky, D. and Martin, J. H. (2018). Speech and language processing (draft). *Chapter 5: Logistic Regression*, 19:2019.
- Jurkowitz, M., Mitchell, A., Shearer, E., and Walker, M. (2020). Us media polarization and the 2020 election: A nation divided. *Pew Research Center*.
- Kaplan, A. M. and Haenlein, M. (2010). Users of the world, unite! the challenges and opportunities of social media. *Business horizons*, 53(1):59–68.

- Kelman, H. C. (1958). Compliance, identification, and internalization three processes of attitude change. *Journal of conflict resolution*, 2(1):51–60.
- Khan, R., Misra, K., and Singh, V. (2013). Ideology and brand consumption. *Psychological Science*, 24(3):326–333.
- Khatua, A., Khatua, A., and Cambria, E. (2020). Predicting political sentiments of voters from twitter in multi-party contexts. *Applied Soft Computing*, 97:106743.
- Kidwell, B., Farmer, A., and Hardesty, D. M. (2013). Getting liberals and conservatives to go green: Political ideology and congruent appeals. *Journal of Consumer Research*, 40(2):350–367.
- King, G., Pan, J., and Roberts, M. E. (2014). Reverse-engineering censorship in china: Randomized experimentation and participant observation. *Science*, 345(6199).
- Kirchgessner, M., Leroy, V., Amer-Yahia, S., and Mishra, S. (2016). Testing interest-iness measures in practice: A large-scale analysis of buying patterns. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 547–556.
- Kosinski, M., Matz, S. C., Gosling, S. D., Popov, V., and Stillwell, D. (2015). Facebook as a research tool for the social sciences: Opportunities, challenges, ethical considerations, and practical guidelines. *American Psychologist*, 70(6):543–556.
- Kosinski, M., Stillwell, D., and Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805.
- Koumchatzky, N. and Andryeyev, A. (2017). Using deep learning at scale in twitter’s timelines.
- Krackhardt, D. and Stern, R. N. (1988). Informal networks and organizational crises: An experimental simulation. *Social psychology quarterly*, pages 123–140.
- Kreiss, D. (2016). Seizing the moment: The presidential campaigns’ use of twitter during the 2012 electoral cycle. *New media & society*, 18(8):1473–1490.
- Kreutz, T. and Daelemans, W. (2018). Enriching general sentiment lexicons for domain-specific polarity classification. *Manuscript in preparation*.
- Kristensen, J. B., Albrechtsen, T., Dahl-Nielsen, E., Jensen, M., Skovrind, M., and Bornakke, T. (2017). Parsimonious data: How a single facebook like predicts voting behavior in multiparty systems. *PloS one*, 12(9):1–12.
- Kunst, A. (2020). Social media activities in belgium 2020. *Statistica*.
- Lai, S., Xu, L., Liu, K., and Zhao, J. (2015). Recurrent convolutional neural networks for text classification. In *AAAI*, volume 333, pages 2267–2273.
- Lazer, D., Pentland, A. S., Adamic, L., Aral, S., Barabasi, A. L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., et al. (2009). Life in the network: the coming age of computational social science. *Science (New York, NY)*, 323(5915):721.
- Lazer, D. M., Pentland, A., Watts, D. J., Aral, S., Athey, S., Contractor, N., Freelon, D., Gonzalez-Bailon, S., King, G., Margetts, H., et al. (2020). Computational social science: Obstacles and opportunities. *Science*, 369(6507):1060–1062.
- LD, D. B. and Raj, E. D. (2017). Flocking based evolutionary computation strategy for measuring centrality of online social networks. *Applied Soft Computing*, 58:495–516.
- Lee, A. H.-Y. (2020). How the politicization of everyday activities affects the public sphere: The effects of partisan stereotypes on cross-cutting interactions. *Political*

- Communication*, pages 1–20.
- Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791.
- Lelkes, Y. (2016). Mass polarization: Manifestations and measurements. *Public Opinion Quarterly*, 80(S1):392–410.
- Lijphart, A. (2012). *Patterns of democracy: Government forms and performance in thirty-six countries*. Yale University Press.
- Lipset, S. M. and Rokkan, S. (1967). *Party systems and voter alignments: cross-national perspectives*. Free press.
- Liu, D. Y., Taylor, C. E., Bridgeman, A. J., Bartimote-Aufflick, K., and Pardo, A. (2016). Empowering instructors through customizable collection and analyses of actionable information. In *PCLA@ LAK*, pages 3–9.
- Lockhart, R., Taylor, J., Tibshirani, R. J., and Tibshirani, R. (2014). A significance test for the lasso. *Annals of statistics*, 42(2):413.
- Luhiste, M., Popa, S. A., Theocharis, Y., Barberá, P., Praet, S., Fazekas, Z., and Tucker, J. (2021). When does Abuse and Harassment Marginalize Female Political Voices on Social Media? *To be submitted*.
- Macafee, T. (2013). Some of these things are not like the others: Examining motivations and political predispositions among political facebook activity. *Computers in Human Behavior*, 29(6):2766–2775.
- Mack, Graig (2020). Social media: Engaging the world in australia’s bushfire crisis.
- MacWilliams, M. C. (2015). Forecasting congressional elections using facebook data. *PS: Political Science & Politics*, 48(4):579–583.
- Magurran, A. E. (2013). *Measuring biological diversity*. John Wiley & Sons.
- Maly, I. (2020). Algorithmic populism and the datafication and gamification of the people by flemish interest in belgium. *Trabalhos em Linguística Aplicada*, 59(1):444–468.
- Marder, B. (2018). Trumped by context collapse: Examination of ‘liking’ political candidates in the presence of audience diversity. *Computers in Human Behavior*, 79:169–180.
- Martens, D. (2021). *Data Science Ethics - Concepts, Techniques and Cautionary Tales*. CLARENDON PRESS OXFORD.
- Martens, D., Provost, F., Clark, J., and de Fortuny, E. J. (2016). Mining massive fine-grained behavior data to improve predictive analytics. *MIS quarterly*, 40(4):869–888.
- Mason, L. (2018a). Ideologues without Issues: The Polarizing Consequences of Ideological Identities. *Public Opinion Quarterly*, 82(S1):866–887.
- Mason, L. (2018b). *Uncivil Agreement: How Politics Became Our Identity*. University of Chicago Press.
- Matz, S. C. and Netzer, O. (2017). Using big data as a window into consumers’ psychology. *Current opinion in behavioral sciences*, 18:7–12.
- McNicholas, P., Murphy, T., and O’Regan, M. (2008). Standardising the lift of an association rule. *Computational Statistics & Data Analysis*, 52(10):4712–4721.
- McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444.

- Medina Serrano, J. C., Papakyriakopoulos, O., and Hegelich, S. (2020). Dancing to the partisan beat: a first analysis of political communication on tiktok. In *12th ACM Conference on Web Science*, pages 257–266.
- Mellon, J. and Prosser, C. (2017). Twitter and facebook are not representative of the general population: Political attitudes and demographics of british social media users. *Research & Politics*, 4(3):2053168017720008.
- Meredith, M. and Ridout, M. (2014). Overview of the overlap package. *R. Proj*, pages 1–9.
- Messing, S., Kessel, V., and Hughes, A. (2017). Sharing the news in a polarized congress. *Technical Report*.
- Messing, S. and Westwood, S. J. (2014). Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online. *Communication research*, 41(8):1042–1063.
- Moeller, J., Helberger, N., et al. (2018). Beyond the filter bubble: Concepts, myths, evidence and issues for future debates.
- Monroe, B. L., Pan, J., Roberts, M. E., Sen, M., and Sinclair, B. (2015). No! formal theory, causal inference, and big data are not contradictory trends in political science. *PS, Political Science & Politics*, 48(1):71.
- Mosleh, M., Pennycook, G., and Rand, D. G. (2020). Self-reported willingness to share political news articles in online surveys correlates with actual sharing on twitter. *Plos one*, 15(2):1–9.
- Müller, K. and Schwarz, C. (2020). From hashtag to hate crime: Twitter and anti-minority sentiment. *Available at SSRN 3149103*.
- Mummolo, J. and Nall, C. (2017). Why partisans do not sort: The constraints on political segregation. *The Journal of Politics*, 79(1):45–59.
- Munger, K., Bonneau, R., Nagler, J., and Tucker, J. A. (2014). Elites tweet to get feet off the streets: Measuring regime response to protest using social media.
- Mutz, D. C. (2006). *Hearing the other side: Deliberative versus participatory democracy*. Cambridge University Press.
- Myllylahti, M. (2018). An attention economy trap? an empirical investigation into four news companies’ facebook traffic and social media revenue. *Journal of Media Business Studies*, 15(4):237–253.
- Nadeau, C. and Bengio, Y. (2000). Inference for the generalization error. In *Advances in neural information processing systems*, pages 307–313.
- Nagler, J. and Tucker, J. A. (2015). Drawing inferences and testing theories with big data. *PS: Political Science & Politics*, 48(1):84–88.
- Narayanan, A. and Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125. IEEE.
- NBC News (2020). A movement, a slogan, a rallying cry: How black lives matter changed america’s view on race.
- Newman, B. J. and Bartels, B. L. (2011). Politics at the checkout line: Explaining political consumerism in the united states. *Political Research Quarterly*, 64(4):803–817.
- Newman, M. (2010). *Networks: An introduction*. Oxford university press.

- Newman, M. E. (2006). Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582.
- Newman, N., Fletcher, R., Kalogeropoulos, A., and Nielsen, R. (2019a). *Reuters institute digital news report 2019*, volume 2019. Reuters Institute for the Study of Journalism.
- Newman, N., Fletcher, R., Kalogeropoulos, A., and Nielsen, R. K. (2019b). Reuters institute digital news report 2019. *Reuters Institute for the Study of Journalism*.
- Nulty, P., Theocharis, Y., Popa, S. A., Parnet, O., and Benoit, K. (2016). Social media and political communication in the 2014 elections to the european parliament. *Electoral studies*, 44:429–444.
- O’donnell, G. and Schmitter, P. C. (2013). *Transitions from authoritarian rule: Tentative conclusions about uncertain democracies*. JHU Press.
- Office for Human Research Protections, U. D. o. H. and Services, H. (2020). International compilation of human research standards.
- O’callaghan, D., Greene, D., Carthy, J., and Cunningham, P. (2015). An analysis of the coherence of descriptors in topic modeling. *Expert Systems with Applications*, 42(13):5645–5657.
- Panzeri, S. and Treves, A. (1996). Analytical estimates of limited sampling biases in different information measures. *Network: Computation in neural systems*, 7(1):87–107.
- Parmelee, J. H. and Bichard, S. L. (2011). *Politics and the Twitter revolution: How tweets influence the relationship between political leaders and the public*. Lexington Books.
- Paul, D., Li, F., Teja, M. K., Yu, X., and Frost, R. (2017). Compass: Spatio temporal sentiment analysis of us election what twitter says! In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1585–1594. ACM.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011a). Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011b). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Peeters, J., Van Aelst, P., and Praet, S. (2019). Party ownership or individual specialization? A comparison of politicians’ individual issue attention across three different agendas. *Party Politics*, pages 1–12.
- Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., and Rand, D. G. (2020). Shifting attention to accuracy can reduce misinformation online. *Nature*.
- Perriam, J., Birkbak, A., and Freeman, A. (2020). Digital methods in a post-api environment. *International Journal of Social Research Methodology*, 23(3):277–290.
- Perrin, A. and Anderson, M. (2019). Share of us adults using social media, including facebook, is mostly unchanged since 2018. *Pew Research Center*, 10.
- Persily, N. (2017). The 2016 us election: Can democracy survive the internet? *Journal of democracy*, 28(2):63–76.
- Petrocik, J. (1989). The theory of issue ownership: Issues, agendas, and electoral coalitions in the 1988 elections. In *annual meeting of the American Political Science*

- Association.*
- Petrocik, J. R. (1996). Issue ownership in presidential elections, with a 1980 case study. *American journal of political science*, pages 825–850.
- Pettigrew, T. F. (1998). Intergroup contact theory. *Annual review of psychology*, 49(1):65–85.
- Pew Research Center (2016). Election 2016: Campaigns as a direct source of news.
- Piazza, A., Zagel, C., Haeske, J., and Bodendorf, F. (2017). Do you like according to your lifestyle? a quantitative analysis of the relation between individual facebook likes and the users lifestyle. In *International Conference on Applied Human Factors and Ergonomics*, pages 128–139. Springer.
- Pizarro, J., Guerrero, E., and Galindo, P. L. (2002). Multiple comparison procedures applied to model selection. *Neurocomputing*, 48(1-4):155–173.
- Praet, S., Daelemans, W., Kreutz, T., Van Aelst, P., Walgrave, S., and Martens, D. (2018a). Issue communication by political parties on Twitter. Data Science, Journalism and Media workshop at KDD Conference, London, UK.
- Praet, S., Guess, A. M., Tucker, J. A., Bonneau, R., and Nagler, J. (2021a). What’s Not to Like? Facebook Page Likes Reveal Limited Polarization in Lifestyle Preferences. *R&R: Political Communication*.
- Praet, S. and Martens, D. (2017). Efficient Parcel Delivery by Predicting Customers’ Locations. IFORS: 21st Conference of the International Federation of Operational Research Societies, Quebec, Canada.
- Praet, S. and Martens, D. (2020). Efficient parcel delivery by predicting customers’ locations. *Decision Sciences*, 51(5):1202–1231.
- Praet, S., Van Aelst, P., Daelemans, W., Walgrave, S., Kreutz, T., , and Martens, D. (2018b). Issue Ownership and Communication by Political Parties on Twitter. EURO: 29th European Conference on Operational Research, Valencia, Spain.
- Praet, S., Van Aelst, P., Daelemans, W., Walgrave, S., Kreutz, T., Peeters, J., and Martens, D. (2020a). Comparing automatic content analysis methods to distinguish issue communication by political parties on Twitter. *Accepted: Computational Communication Research*.
- Praet, S., Van Aelst, P., and Martens, D. (2020b). Partisanship in everyday life: predictive modeling with Facebook like data. ORBEL34: 34th conference of the Belgian Operational Research Society, Lille, France.
- Praet, S., Van Aelst, P., and Martens, D. (2021b). Patterns of democracy? Social network analysis of parliamentary Twitter networks in 12 countries. *R&R: Online Social Networks and Media*.
- Praet, S., Van Aelst, P., Van der Veekend, S., and Martens, D. (2019a). Predictive modeling on Facebook likes to gain insights in voter profiles in a multi-party system. 5th International Conference on Computational Social Science IC2S2, Amsterdam, The Netherlands.
- Praet, S., Van Aelst, P., Van der Veekend, S., and Martens, D. (2019b). What I Do versus Who I Am. Comparing Facebook likes with survey data to gain insights in political preference. EURO: 30th European Conference on Operational Research, Dublin, Ireland.

- Praet, S., Van Aelst, P., Van erkel, P., Van der Veeken, S., and Martens, D. (2020c). How Facebook likes can add lifestyle politics to complement traditional election survey research. *R&R: EPJ Data Science*.
- Prior, M. (2009). Improving media effects research through better measurement of news exposure. *The Journal of Politics*, 71(3):893–908.
- Prior, M. (2013). The challenge of measuring media exposure: Reply to dilliplane, goldman, and mutz. *Political Communication*, 30(4):620–634.
- Provost, F. and Fawcett, T. (2013). *Data Science for Business: What you need to know about data mining and data-analytic thinking*. " O'Reilly Media, Inc."
- Provost, F., Martens, D., and Murray, A. (2015). Finding similar mobile consumers with a privacy-friendly geosocial design. *Information Systems Research*, 26(2):243–265.
- Purhonen, S. and Heikkilä, R. (2017). Food, music and politics: The interweaving of culinary taste patterns, highbrow musical taste and conservative attitudes in finland. *Social Science Information*, 56(1):74–97.
- Quattrociocchi, W., Scala, A., and Sunstein, C. R. (2016). Echo chambers on facebook. *Available at SSRN 2795110*.
- Quote Investigator (2021). To err is human; to really foul things up requires a computer.
- Raats, T., Picone, I., Van Aelst, P., and Paulussen, S. (2021). De onpartijdigheid van het VRT-aanbod. Technical report, Vrije Universiteit Brussel en Universiteit Antwerpen.
- Radecki, M. (2016). Party leadership in poland in comparative perspective. *Annales Universitatis Mariae Curie-Skłodowska, sectio K–Politologia*, 22(1):135.
- Ramon, Y., Martens, D., Evgeniou, T., and Praet, S. (2021). Metafeatures-based rule-extraction for classifiers on behavioral and textual data. *Accepted: Machine Learning*.
- Rathi, R. (2019). Effect of cambridge analytica's facebook ads on the 2016 us presidential election. *Towards Data Science*.
- Reardon, M. (2021). The Biden presidency: What it means for tech.
- Reynolds, A., Reilly, B., Council, N. R., et al. (1999). *Electoral systems and conflict in divided societies*. National Academies Press.
- Robles, J. M., Velez, D., De Marco, S., Rodríguez, J. T., and Gomez, D. (2020). Affective homogeneity in the spanish general election debate. a comparative analysis of social networks political agents. *Information, Communication & Society*, 23(2):216–233.
- Roos, J. M. and Shachar, R. (2014). When kerry met sally: Politics and perceptions in the demand for movies. *Management Science*, 60(7):1617–1631.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215.
- Saunders, K. L. and Abramowitz, A. I. (2004). Ideological realignment and active partisans in the american electorate. *American Politics Research*, 32(3):285–309.
- Schat, E., van de Schoot, R., Kouw, W. M., Veen, D., and Mendrik, A. M. (2020). The data representativeness criterion: Predicting the performance of supervised

- classification based on data set similarity. *Plos one*, 15(8):e0237009.
- Scherpereel, J. A., Wohlgemuth, J., and Schmelzinger, M. (2017). The adoption and use of twitter as a representational tool among members of the european parliament. *European politics and society*, 18(2):111–127.
- Schmidt, A. and Wiegand, M. (2017). A survey on hate speech detection using natural language processing. In *Proceedings of the fifth international workshop on natural language processing for social media*, pages 1–10.
- Schober, M. F., Pasek, J., Guggenheim, L., Lampe, C., and Conrad, F. G. (2016). Social media analyses for social measurement. *Public opinion quarterly*, 80(1):180–211.
- Schwartz, D. G., Yahav, I., and Silverman, G. (2017). News censorship in online social networks: A study of circumvention in the commentsphere. *Journal of the Association for Information Science and Technology*, 68(3):569–582.
- Scott, D. W. (2015). *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons.
- Seaver, N. (2019). Knowing algorithms. *Digital STS*, pages 412–422.
- Selway, J. S. (2011). The measurement of cross-cutting cleavages and other multidimensional cleavage structures. *Political Analysis*, 19(1):48–65.
- Settle, J. E. (2018). *Frenemies: How Social Media Polarizes America*. Cambridge University Press.
- Settle, J. E., Bond, R. M., Coviello, L., Fariss, C. J., Fowler, J. H., and Jones, J. J. (2016). From posting to voting: The effects of political competition on online political engagement. *Political Science Research and Methods*, 4(2):361–378.
- Sevenans, J., Albaugh, Q., Shahaf, T., Soroka, S., and Walgrave, S. (2014). The automated coding of policy agendas: A dictionary based approach (v. 2.0.). In *7th annual Comparative Agendas Project (CAP) conference, Konstanz, June*, pages 12–14.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423.
- Shi, J., Lai, K. K., Hu, P., and Chen, G. (2017a). Understanding and predicting individual retweeting behavior: Receiver perspectives. *Applied Soft Computing*, 60:844–857.
- Shi, Y., Mast, K., Weber, I., Kellum, A., and Macy, M. (2017b). Cultural fault lines and political polarization. In *Proceedings of the 2017 ACM on Web Science Conference*, pages 213–217.
- Shmueli, G. (2017). Analyzing behavioral big data: methodological, practical, ethical, and moral issues. *Quality Engineering*, 29(1):57–74.
- Shmueli, G. et al. (2010). To explain or to predict? *Statistical science*, 25(3):289–310.
- Shu, K., Sliva, A., Wang, S., Tang, J., and Liu, H. (2017). Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*, 19(1):22–36.
- Siegel, A., Nikitin, E., Barberá, P., Sterling, J., Pullen, B., Bonneau, R., Nagler, J., and Tucker, J. A. (2019). Trumping hate on twitter? online hate in the 2016 us election and its aftermath. *Social Media and Political Participation, New York University*.
- Siegel, A. A. and Badaan, V. (2020). #nozsectarianism: Experimental approaches to reducing sectarian hate speech online. *American Political Science Review*, 114(3):837–855.



- Siegel, D. A. (2011). Social networks in comparative perspective. *PS: Political Science and Politics*, 44(1):51–54.
- Sikder, O., Smith, R. E., Vivo, P., and Livan, G. (2020). A minimalistic model of bias, polarization and misinformation in social networks. *Scientific reports*, 10(1):1–11.
- Singel, R. (2009). Netflix spilled your brokeback mountain secret, lawsuit claims. *Threat Level (blog)*, *Wired*.
- Sîrbu, A., Pedreschi, D., Giannotti, F., and Kertész, J. (2019). Algorithmic bias amplifies opinion fragmentation and polarization: A bounded confidence model. *PloS one*, 14(3):e0213246.
- Solove, D. J. (2005). A taxonomy of privacy. *U. Pa. L. Rev.*, 154:477.
- Sorzano, C. O. S., Vargas, J., and Montano, A. P. (2014). A survey of dimensionality reduction techniques. *arXiv preprint arXiv:1403.2877*.
- Sperandei, S. (2014). Understanding logistic regression analysis. *Biochemia medica*, 24(1):12–8.
- Stankova, M., Praet, S., Martens, D., and Provost, F. (2021). Node classification over bipartite graphs through projection. *Machine Learning*, 110:37–87.
- StatBel (2018). Onderwijsniveau van de belgische bevolking van 15-64 jaar. <https://bestat.statbel.fgov.be/bestat/crosstable.xhtml?view=631b4535-7a63-4695-967f-fe42238ee9af>.
- StatCounter (2019). Social media stats 2019.
- Statista (2020). Facebook user distribution in belgium 2020, by age group and gender.
- Stevens, K., Kegelmeyer, P., Andrzejewski, D., and Buttler, D. (2012). Exploring topic coherence over many models and many topics. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 952–961. Association for Computational Linguistics.
- Stier, S., Breuer, J., Siegers, P., and Thorson, K. (2019). Integrating survey data and digital trace data: Key issues in developing an emerging field. *Social Science Computer Review*, 38(5):503–516.
- Stier, S., Posch, L., Bleier, A., and Strohmaier, M. (2017). When populists become popular: comparing facebook use by the right-wing movement pegida and german political parties. *Information, Communication & Society*, 20(9):1365–1388.
- Stukal, D., Sanovich, S., Bonneau, R., and Tucker, J. A. (2017). Detecting bots on russian political twitter. *Big data*, 5(4):310–324.
- Sydell, L. (2016). We tracked down a fake-news creator in the suburbs. here’s what we learned. *National Public Radio*, 23.
- Tangherlini, T. R., Shahsavari, S., Shahbazi, B., Ebrahimzadeh, E., and Roychowdhury, V. (2020). An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, pizzagate and storytelling on the web. *PloS one*, 15(6):e0233879.
- Teernstra, L., Uitermark, J., and Tornberg, P. (2018). Politicians’ retweet networks: Political coalitions and divisions on twitter in 23 countries. Working paper.
- Terechshenko, Z., Linder, F., Padmakumar, V., Liu, F., Nagler, J., Tucker, J. A., and Bonneau, R. (2020). A comparison of methods in political science text classification: Transfer learning language models for politics. In *2020 Annual Meeting of the Society*

- for *Political Methodology*.
- Thomassen, J. and Schmitt, H. (1997). Policy representation. *European Journal of Political Research*, 32(2):165–184.
- Tibshirani, R., Wainwright, M., and Hastie, T. (2015). *Statistical learning with sparsity: the lasso and generalizations*. Chapman and Hall/CRC.
- Time (2010). Person of the year 2010.
- Tobback, E. and Martens, D. (2019). Retail credit scoring using fine-grained payment data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182(4):1227–1246.
- Toff, B. (N.d.). Prime time television in red and blue: Partisan divides in entertainment media. *Working Paper*.
- Tornes, A. and Trujillo, L. (2021). Enabling the future of academic research with the twitter api.
- Townsend, L. and Wallace, C. (2016). Social media research: A guide to ethics. *University of Aberdeen*, 1:16.
- Tresch, A., Lefevere, J., and Walgrave, S. (2017). How parties’ issue emphasis strategies vary across communication channels: The 2009 regional election campaign in belgium. *Acta Politica*, pages 1–23.
- Tucker, J. A., Theocharis, Y., Roberts, M. E., and Barberá, P. (2017). From liberation to turmoil: Social media and democracy. *Journal of democracy*, 28(4):46–59.
- Tufféry, S. (2011). *Data mining and statistics for decision making*. Wiley Chichester.
- Tworzecki, H. (2019). Poland: a case of top-down polarization. *The ANNALS of the American Academy of Political and Social Science*, 681(1):97–119.
- UN General Assembly (1948). Universal declaration of human rights. *UN General Assembly*, 302(2).
- Urman, A. (2017). *Opposition in Putin’s Russia: Resisting Propaganda*. PhD thesis, Central European University.
- Urman, A. (2020). Context matters: political polarization on twitter from a comparative perspective. *Media, culture & society*, 42(6):857–879.
- Vaccari, C., Valeriani, A., Barberá, P., Jost, J. T., Nagler, J., and Tucker, J. A. (2016). Of echo chambers and contrarian clubs: Exposure to political disagreement among german and italian users of twitter. *Social media+ society*, 2(3):2056305116664221.
- Valente, T. W. (1996). Social network thresholds in the diffusion of innovations. *Social networks*, 18(1):69–89.
- Van Dalen, A., Fazekas, Z., Klemmensen, R., and Hansen, K. M. (2015). Policy considerations on facebook: Agendas, coherence, and communication patterns in the 2011 danish parliamentary elections. *Journal of Information Technology & Politics*, 12(3):303–324.
- Van Dijck, J. and Poell, T. (2013). Understanding social media logic. *Media and communication*, 1(1):2–14.
- Van Ditmars, M. M., Maggini, N., and van Spanje, J. (2020). Small winners and big losers: strategic party behaviour in the 2017 dutch general election. *West European Politics*, 43(3):543–564.
- Van Engelen, J. E. and Hoos, H. H. (2020). A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440.

- Van Erkel, P., Thijssen, P., and Van Aelst, P. (2014). Vier ideologische dimensies, één breuklijn. *Stichting Gerrit Kreveld: Samenleving en Politiek*, pages 1–12.
- van Kessel, P., Widjaya, R., Shah, S., Smith, A., and Hughes, A. (2020). The congressional social media landscape.
- Van Santen, R., Helfer, L., and Van Aelst, P. (2015). When politics becomes news: An analysis of parliamentary questions and press coverage in three west european countries. *Acta Politica*, 50(1):45–63.
- Van Vliet, L., Törnberg, P., and Uitermark, J. (2020). The twitter parliamentary database: Analyzing twitter politics across 26 countries. *PloS one*, 15(9):e0237073.
- Vanhoeyveld, J., Martens, D., and Peeters, B. (2020). Value-added tax fraud detection with scalable anomaly detection techniques. *Applied Soft Computing*, 86:105895.
- Vanhoeyveld, J., Martens, D., Peeters, B., Van de Velde, S., Van Eekert, N., Van Assche, K., Wouters, E., Verschueren, H., Vansweevelt, T., Bubalo, L., et al. (2019). Customs fraud detection: Assessing the value of behavioural and high-cardinality data under the imbalanced learning issue. *Pattern analysis and applications, ISSN*, pages 1433–7541.
- Vargo, C. J., Guo, L., McCombs, M., and Shaw, D. L. (2014). Network issue agendas on twitter during the 2012 us presidential election. *Journal of Communication*, 64(2):296–316.
- Vázquez, E. G., Escolano, A. Y., Riaño, P. G., and Junquera, J. P. (2001). Repeated measures multiple comparison procedures applied to model selection in neural networks. In *International Work-Conference on Artificial Neural Networks*, pages 88–95. Springer.
- Vera, E. R. and Schupp, T. (2006). Network analysis in comparative social sciences. *Comparative Education*, 42(3):405–429.
- Verbeke, W., Martens, D., and Baesens, B. (2014). Social network analysis for customer churn prediction. *Applied Soft Computing*, 14:431–446.
- Volkens, A., Burst, T., Krause, W., Lehmann, P., Matthiess, T., Merz, N., Regel, S., Wessels, B., and Zehnter, L. (2020). The manifesto data collection. manifesto project (mrg/cmp/marpor). version 2020a.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Wagner, M. (2020). Affective polarization in multiparty systems. *Electoral Studies*, page 102199.
- Wagner, M. and Meyer, T. M. (2014). Which issues do parties emphasise? salience strategies and party organisation in multiparty systems. *West European Politics*, 37(5):1019–1045.
- Walgrave, S., Tresch, A., and Lefevere, J. (2015). The conceptualisation and measurement of issue ownership. *West European Politics*, 38(4):778–796.
- Wallach, H. (2018). Computational social science  $\neq$  computer science+ social data. *Communications of the ACM*, 61(3):42–44.
- Wang, C. and Blei, D. M. (2011). Collaborative topic modeling for recommending scientific articles. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 448–456.

- We Are Social and Hootsuite (2020). Digital 2020: Global digital overview. *Datareportal*.
- Wells, C. and Thorson, K. (2017). Combining big data and survey techniques to model effects of political content flows in facebook. *Social Science Computer Review*, 35(1):33–52.
- Westin, A. F. (1968). Privacy and freedom. *Washington and Lee Law Review*, 25(1):166.
- Westwood, S. J., Iyengar, S., Walgrave, S., Leonisio, R., Miller, L., and Strijbis, O. (2018). The tie that divides: Cross-national evidence of the primacy of partyism. *European Journal of Political Research*, 57(2):333–354.
- Wojcik, S. and Hughes, A. (2019). Sizing up twitter users. *Pew Research Center*.
- Worldbank (2019). Population, total.
- Wu, L., Morstatter, F., and Liu, H. (2018). Slangs: building, expanding and using a sentiment dictionary of slang words for short-text sentiment classification. *Language Resources and Evaluation*, 52(3):839–852.
- Yamaya, S., Bhadani, S., Flammioni, A., Menczer, F., Nyhan, B., and Ciampaglia, G. L. (2020). Political audience diversity and news quality.
- Young, D. G. (2019). *Irony and outrage: The polarized landscape of rage, fear, and laughter in the United States*. Oxford University Press, USA.
- Zilinsky, J., Vaccari, C., Nagler, J., and Tucker, J. A. (2020). Don’t republicans tweet too? using twitter to assess the consequences of political endorsements by celebrities. *Perspectives on Politics*, 18(1):144–160.
- Zimmer, M. (2010). But the data is already public: on the ethics of research in facebook. *Ethics and information technology*, 12(4):313–325.
- Zirn, C., Glavaš, G., Nanni, F., Eichorts, J., and Stuckenschmidt, H. (2016). Classifying topics and detecting topic shifts in political manifestos.
- Zook, M., Barocas, S., Boyd, D., Crawford, K., Keller, E., Gangadharan, S. P., Goodman, A., Hollander, R., Koenig, B. A., Metcalf, J., et al. (2017). Ten simple rules for responsible big data research. *PLoS Computational Biology*, 13(3):1–10.

## Part IV

## APPENDICES



## Appendices

---

### 9.1 Issue communication on Twitter

#### 9.1.1 *Evaluation of the CAP dictionary per issue*

Table 9.1: Intercoder reliability (Cohen's Kappa), per issue for the annotated tweets.

Issue	ICR
Macroeconomics	0.61
Human rights	0.43
Health	0.52
Agriculture	0.79
Labor and employment	0.60
Immigration	0.69
Education	0.69
Environment	0.68
Energy	0.73
Transportation	0.64
Law and crime	0.47
Social welfare	0.31
Community development	0.30
Banking and finance	0.56
Defense	0.51
Science and technology	0.18
Foreign trade	0.00
International affairs	0.40
Government operations	0.41
Public lands and water	0.00
Culture and arts	0.41

The issues with low (to zero) ICR have low (to no) occurrence.

Table 9.2: Evaluation of the CAP dictionary in terms of precision, recall and F1 score per issue.

Issue	Precision	Recall	F1 score
Macroeconomics	43%	13%	20%
Human rights	44%	19%	27%
Health	63%	21%	31%
Agriculture	66%	20%	31%
Labor and employment	54%	28%	37%
Immigration	74%	22%	34%
Education	72%	36%	48%
Environment	65%	35%	45%
Energy	66%	36%	47%
Transportation	77%	48%	59%
Law and crime	61%	16%	25%
Social welfare	66%	4%	8%
Community development	29%	12%	17%
Banking and finance	60%	8%	14%
Defense	69%	13%	21%
Science and technology	50%	1%	2%
Foreign trade	0%	0%	0%
International affairs	27%	7%	11%
Government operations	46%	3%	5%
Public lands and water	0%	0%	0%
Culture and arts	60%	5%	9%



Table 9.3: Number and percentage of tweets that was assigned a certain CAP issue using the CAP dictionary.

Issue	Number of tweets	Percentage of tweets
Macroeconomics	911	2%
Human rights	517	1%
Health	780	1%
Agriculture	393	1%
Labor and employment	1265	2%
Immigration	1170	2%
Education	1772	3%
Environment	767	1%
Energy	598	1%
Transportation	2277	4%
Law and crime	1138	2%
Social welfare	521	1%
Community development	277	0%
Banking and finance	256	0%
Defense	291	1%
Science and technology	55	0%
Foreign trade	104	0%
International affairs	735	1%
Government operations	454	1%
Public lands and water	17	0%
Culture and arts	116	0%
No issue	43245	76%

### 9.1.2 *Extended dictionary with word embedding*

The dictionary maps keywords to their respective political issues and aims to be very precise, with keywords having a very distinct meaning and low probability to be present in one of the other issues. For analysis of short social media texts such as tweets, in which very few words are present, this precision is less important and coverage with the expert dictionary is of more concern. To extend the indicator words in the original dictionary, we use word embeddings trained on a large corpus of political social media data (Kreutz and Daelemans, 2018). The word embeddings encode a numerical vector per word, which contains the point-wise mutual information (PMI) with other words in the corpus. Using these vectors, we can find candidate words that are semantically similar to the keywords already present in the dictionary, using a cosine-similarity of 0.6 or higher. The candidates were then manually inspected and filtered to contain only words that extend coverage of the expert issues without clearly impairing their delineation. Using word embeddings in this way, we were able to extend the keywords from an average of 87 per expert issue to 157 per expert issue and consequently, 85% of the tweets could be assigned at least one issue (compared to only 24 % for the original dictionary).

The extended dictionary was also tested on a random subset of 9,280 tweets that was manually coded for the 21 CAP issues. Accuracy of the extended dictionary is 20%, recall is 35% and precision is 39%. Although recall and coverage could be increased by extending the dictionary, the precision is much lower than that of the original dictionary. For this reason we decided to apply the original dictionary in this research, despite the low coverage. This shows that accurately extending the existing dictionaries is still a difficult challenge.

9.1.3 *Most discriminative features for the data-driven representations*

Table 9.4: The most discriminative features when using the BoW approach and their three most related CAP issues. Named entities are printed in capital letters.

Party	Most discriminative features	Corresponding CAP issues
Groen	itcanbedifferent, greenworks, lowerhouse, glyphosate, meyremalmaci, climate ambition, morehealthy, changecongress, antwerpcondoit, advertising, widening change, concrete stop, cyclist, green, screening, whistle blower, forest, air pollution, pesticide, climate generation, c'est (French), youthforclimate, climate top, incomprehensible, air, fairer, hormone disruptor, glyphosate, takeNUMBER, hood, practical test, longlivepolitics, e.g., climate policy, serious, audit, kristofcalvo, position, flemish, terzakets, majority, complete, unworthy, unbelievable, poverty line	1. Environment 2. / 3. /
sp.a	security of care, wetakecare, municipal works, schaarbeek, bredene, securityforall, stopthedebtindustry, flemish government, goleft, security, johncrombez, debt industry, assets, molenbeek, proposal, nuclear weapon, new battle, weapon embargo, fail, deposit (for packaging), future budget, beach, care crisis, plastic, resolution, veviba (meat company), replacement, nuclear, youarewhatyoueat, sp.a, history of bredene, throwback, contest, vanovertveldt, voted out, water bill, litter, saving, reynders, feed, plow, reading, profit, crazy, weak	1. Social welfare 2. Environment 3. Macroeconomics
CD&V	thewayforward, quality of life, socialeurope, gtgen, justice, wbeke, bike is king, social, climate court, cd&v, traffic jam idea, safe traffic, consultation, peeterskrisNUMBER, residential care centers, teacher, jokeschauvliege, mobility budget, info, koengeensNUMBER, belgians, homeinthecity, crevils, improve, tooth, renew, brexit, worker, inheritance law, social right, electrical, movingsafely, callNUMBER, economic, elderly care, belgiangovernment, simpler, school construction, climate pact, quitting principle, close to you, recommendation, opening, school year, servaisv	1. Social welfare 2. Transportation 3. Education
Open Vld	justdoit, positiveforward, vilvoorde, must (ENG), pedestrian son, liberal, liberal, Sint-truiden, basic income, etc., europe, reform, ambitious, plenary, children (ENG), united (ENG), read, lost, survive, miscellaneous, facebook, subway, would (ENG), unsupported, entrepreneur, agriculture, proud (ENG), think (ENG), could (ENG), need (ENG), dry, closer, iameuropean, city hall, right (ENG), unity (ENG), futureofeurope (ENG), humanright (ENG), minor, Brussels, entrepreneur, strategy (ENG), weareurope (ENG), speech (ENG), task	1. International affairs 2. Macroeconomics 3. Banking and finance
NVA	member of parliament, good news, pride heritage, prisoner, herental, left, vdag, marrakesh coalition, meanwhile, heritage, lgbthistorymonth, member of the European parliament, works of change, minority government, animal welfare, tg, flanders, prime minister, budgetNUMBER, self-determination, rajoy, flemishNUMBER, civil service, marrakesh coalition, structure, homeland, policeman, transit migration, change, via, migrant, gene, union, factor, restriction, catalan, repression, hear, yourpowerfulmanagement, say, excellent, steenokkerzeel, restoration, maybe, prosperity	1. Immigration 2. Government operations 3. Law and crime
VB	immigration, tomvangriek, islamization, vlparl, immigration pact, mass immigration, islam, alien, immigration stop, immigrant, mosque, cordon, mosque, community, population, illegal, immigration policy, asylum seeker, multicultural, border, flandersoursagain, concerning, URL, real, scum, immigrant, cause, country, people, people, terrorist, stop immigration, liberty, independence, ourpeoplefirst, protect our people, muslim, headscarf, so-called, government, even, elite, pact, madness	1. Immigration 2. Government operations 3. /

Table 9.5: The most discriminative features when using the topic modeling representation and their most related CAP issues. Named entities are printed in capital letters.

Party	Most discriminative features	Corresponding CAP issues
Groen	<ol style="list-style-type: none"> <li>1. itcanbedifferent, green, deochtend (radio program), climate, air, work green, meyremalmaci, lower house, kristofcalco, your, poverty, wouterdevriendt, climate policy, plan, honest</li> <li>2. WORK_OF_ART, NUMBERday, flemish parliament, so, vtmnieuws, tomvangriek, antwerpcanoit, zaak, koengeensNUMBER, according to, petermerteng, wbeke, youthforclimate, URL, get</li> <li>3. incomprehensible, guess, advertisement, even, a lot, muyters, only, soil, online, mother tongue, abuse, flight, unacceptable, rent deposit, just</li> </ol>	<ol style="list-style-type: none"> <li>1. Environment</li> <li>2. /</li> <li>3. /</li> </ol>
sp.a	<ol style="list-style-type: none"> <li>1. MENTION, URL, deposit (for packaging), strong, colleague, among others, rightfully, gasses, hearing, deochtend (radio program), proposal, later, success, member of parliament, tonight</li> <li>2. government, federal, parliament, decision, follow, fall, decided, run, flemish, previous, next, prime minister, on behalf of, opposition</li> <li>3. care, for, affordable, security of care, wellbeing, qualitative, quality of life, elderly, informal care, person, elderly care, retirement home, qualitatively, support, quality</li> </ol>	<ol style="list-style-type: none"> <li>1. Environment</li> <li>2. Government operations</li> <li>3. Social welfare</li> </ol>
CD&V	<ol style="list-style-type: none"> <li>1. thewayforward, quality of life, care, thanks to, municipality, job, air quality, plenty, bike is king, reformation, neighbourhood, mobility budget, further, healthy, ambitious</li> <li>2. important, put, step, further, forwards, step, busy, role, direction, because, again, shoulder, follow, measurement, look</li> <li>3. information, URL, from, discuss, during, free, school year, website, from now on, to, dual, number/grade, correct, subscribe, learn</li> </ol>	<ol style="list-style-type: none"> <li>1. Environment</li> <li>2. /</li> <li>3. Education</li> </ol>
Open Vld	<ol style="list-style-type: none"> <li>1. europe, need, new, peopl, social, future, work, today, right, together, must, world, maak, fight, meeting (all in English)</li> <li>2. MENTION, URL, deposit (for packaging), strong, colleague, among others, rightfully, gasses, hearing, deochtend (radio program), proposal, later, success, member of parliament, tonight</li> <li>3. PERSON, URL, prime minister, plus, important (ENG), right (ENG), brussels, migration (ENG), conversation, must, police, ORGANIZATION, question (ENG), us (ENG), one (ENG)</li> </ol>	<ol style="list-style-type: none"> <li>1. International affairs</li> <li>2. Environment</li> <li>3. Immigration</li> </ol>
NVA	<ol style="list-style-type: none"> <li>1. NATIONALITY, URL, meeting, captured, economy, according to, member of parliament, speak, president, colleague, level, political, citizen, violence, nationalities</li> <li>2. say, member of parliament, dare, when, come on, no, enough, debt, alone, nuclear plant, little, money, MENTION, often, opinion</li> <li>3. via, URL, MENTION, member of parliament, representative, save, fiscal, migrant, sail, money, finance, asylum seeker, information, free, security of care</li> </ol>	<ol style="list-style-type: none"> <li>1. Immigration</li> <li>2. Energy</li> <li>3. Immigration</li> </ol>
VB	<ol style="list-style-type: none"> <li>1. ULR, action, and, due to, youngsters, again, care, ready, draw, petition, live, part, right, thanks to, help</li> <li>2. country, border, safe, criminal, population, origin, illegal, deportation, alien, greatest, when, migrant, deportation, hard, nationality</li> <li>3. our, society, protect, safety, propose, economy, society, values, prosperity, and, earn, pride, norm, farmer, resolut</li> </ol>	<ol style="list-style-type: none"> <li>1. Human rights</li> <li>2. Immigration</li> <li>3. Social welfare</li> </ol>

## 9.2 Parliamentary Twitter networks

### 9.2.1 Data collection

The sources that were consulted to collect the names of all members of parliament (Chamber of Representatives and Senate), the president and members of cabinet (Prime minister, Ministers, Secretaries) are provided in Table 9.6. The websites were consulted in May 2018. The total number of politicians is the sum of the members of the parliament (upper house and lower house) and cabinet. Numbers between brackets indicate that there is an overlap between the politicians of this category and other categories. For example, in Poland cabinet members are selected among parliamentarians. Only 17 cabinet members are not part of the parliament and should be added to the total number of politicians. In Belgium we included members of the Federal, Flemish and Walloon parliament and government. The Senate largely consists of members of these parliaments.

Table 9.6: Number of politicians and sources per country.

Country	Category	Number	URL
Belgium	Upper house	(60)	<a href="https://www.senate.be/">https://www.senate.be/</a>
	Lower house	349	<a href="https://www.belgium.be/nl/over_belgie/overheid">https://www.belgium.be/nl/over_belgie/overheid</a>
	Cabinet	34	<a href="https://www.belgium.be/nl/over_belgie/overheid">https://www.belgium.be/nl/over_belgie/overheid</a>
France	Upper house	348	<a href="https://www.senat.fr/senateurs/ump">https://www.senat.fr/senateurs/ump</a>
	Lower house	577	<a href="https://www.assemblee-nationale.fr/dyn/vos-deputes">https://www.assemblee-nationale.fr/dyn/vos-deputes</a>
	Cabinet	37	<a href="https://en.wikipedia.org/wiki/Second_Philippe_government">https://en.wikipedia.org/wiki/Second_Philippe_government</a>
Germany	Upper house	69	<a href="https://www.bundesrat.de/DE/bundesrat/mitglieder/mitglieder-node.html">https://www.bundesrat.de/DE/bundesrat/mitglieder/mitglieder-node.html</a>
	Lower house	709	<a href="https://www.bundestag.de/abgeordnete">https://www.bundestag.de/abgeordnete</a>
	Cabinet	16	<a href="https://www.bundesregierung.de/breg-en/federal-government/cabinet">https://www.bundesregierung.de/breg-en/federal-government/cabinet</a>
Italy	Upper house	321	<a href="https://parlamento17.openpolis.it/lista-dei-parlamentari-in-carica/senato/nome/asc">https://parlamento17.openpolis.it/lista-dei-parlamentari-in-carica/senato/nome/asc</a>
	Lower house	630	<a href="https://parlamento17.openpolis.it/lista-dei-parlamentari-in-carica/camera/nome/asc">https://parlamento17.openpolis.it/lista-dei-parlamentari-in-carica/camera/nome/asc</a>
	Cabinet	20	<a href="https://en.wikipedia.org/wiki/Gentiloni_Cabinet">https://en.wikipedia.org/wiki/Gentiloni_Cabinet</a>
Netherlands	Upper house	75	<a href="https://www.eerstekamer.nl/alle_leden">https://www.eerstekamer.nl/alle_leden</a>
	Lower house	150	<a href="https://www.tweedekamer.nl/kamerleden_en_commissies/alle_kamerleden">https://www.tweedekamer.nl/kamerleden_en_commissies/alle_kamerleden</a>
	Cabinet	27	<a href="https://www.parlement.com/id/vkidc8m3plsz/kabinet_rutte_iii_2017">https://www.parlement.com/id/vkidc8m3plsz/kabinet_rutte_iii_2017</a>
Poland	Upper house	100	<a href="https://www.senat.gov.pl/en/senators/lista-senatorow/">https://www.senat.gov.pl/en/senators/lista-senatorow/</a>
	Lower house	460	<a href="https://www.sejm.gov.pl/poslowie/lista6.htm">https://www.sejm.gov.pl/poslowie/lista6.htm</a>
	Cabinet	17 (32)	<a href="https://en.wikipedia.org/wiki/First_Cabinet_of_Mateusz_Morawiecki">https://en.wikipedia.org/wiki/First_Cabinet_of_Mateusz_Morawiecki</a>
Romania	Upper house	136	<a href="https://www.senat.ro/">https://www.senat.ro/</a>
	Lower house	330	<a href="http://www.cdep.ro/">http://www.cdep.ro/</a>
	Cabinet	21	<a href="http://gov.ro/en/government/the-cabinet-of-ministers/">http://gov.ro/en/government/the-cabinet-of-ministers/</a>
Russia	Upper house	170	<a href="http://www.council.gov.ru/en/structure/members/">http://www.council.gov.ru/en/structure/members/</a>
	Lower house	450	<a href="http://duma.gov.ru/en/duma/deputies/">http://duma.gov.ru/en/duma/deputies/</a>
	Cabinet	34	<a href="https://en.wikipedia.org/wiki/Government_of_Russia">https://en.wikipedia.org/wiki/Government_of_Russia</a>
Spain	Upper house	265	<a href="http://www.senado.es/">http://www.senado.es/</a>
	Lower house	350	<a href="http://www.congreso.es">http://www.congreso.es</a>
	Cabinet	15	<a href="https://en.wikipedia.org/wiki/Government_of_Spain">https://en.wikipedia.org/wiki/Government_of_Spain</a>
UK	Upper house	790	<a href="https://www.parliament.uk/mps-lords-and-offices/lords/">https://www.parliament.uk/mps-lords-and-offices/lords/</a>
	Lower house	650	<a href="https://members.parliament.uk/members/commons">https://members.parliament.uk/members/commons</a>
	Cabinet	(29)	<a href="https://en.wikipedia.org/wiki/Cabinet_of_the_United_Kingdom">https://en.wikipedia.org/wiki/Cabinet_of_the_United_Kingdom</a>
Ukraine	Parliament	421	<a href="https://itd.rada.gov.ua/hrtranslate/Structure/MPs">https://itd.rada.gov.ua/hrtranslate/Structure/MPs</a>
	Cabinet	28	<a href="https://en.wikipedia.org/wiki/Government_of_Ukraine">https://en.wikipedia.org/wiki/Government_of_Ukraine</a>
US	Upper house	100	<a href="https://www.senate.gov/senators/index.htm">https://www.senate.gov/senators/index.htm</a>
	Lower house	435	<a href="https://history.house.gov/Congressional-Overview/Profiles/115th/">https://history.house.gov/Congressional-Overview/Profiles/115th/</a>
	Cabinet	23	<a href="https://en.wikipedia.org/wiki/Cabinet_of_the_United_States">https://en.wikipedia.org/wiki/Cabinet_of_the_United_States</a>

9.2.1.1 *Coding instructions*

For each country, two independent coders with knowledge of the language and political context in the country were asked to manually check the Twitter handles of each politician. Where the two coders did not agree on the correct Twitter handle, the Twitter handle of the politician was inspected by the authors. The coders were provided with the following instructions:

**You will look for the Twitter accounts of**

- Members of parliament (Chamber of Representatives and Senate)
- President and members of cabinet (Prime minister, Ministers, Secretaries)
- Political parties (with seats in parliament)

**What you need to do**

1. Check in the Excel sheet “Overview” whether you believe the links are trustworthy and up-to-date. If not, let me know and look for an up-to-date list of politicians if you can.
2. Check or complete the time period (i.e. the period before the next election)
3. Check in your country sheet whether you find all of the following categories (you can use the filter function on the excel column Category):
  - a) Lower house (if bicameral, otherwise only “parliament”)
  - b) Upper house (if bicameral, otherwise only “parliament”)
  - c) President (if applicable)
  - d) Minister
  - e) Party
4. Check whether all political parties with seats in the parliament are represented. (You can filter the column Category on “Party”)
5. Go through the Excel file name by name and search the name on Twitter to find the correct Twitter account, belonging to this politician. Several users might have the same name, we need the politician’s Twitter account, so read the bibliography carefully. Also, “fake” accounts exists so look for verified

accounts as much as possible. If you do not find the politician's name on Twitter try to check different spellings of the name.

6. If you found the correct Twitter account, write down the following information in the excel sheet of your country:
  - a) Twittername: everything that comes after the @
  - b) Verified: is this a verified account yes/no
  - c) Followers: the amount of followers
  - d) Remarks: if you have remarks (eg. you are not sure about the account, you have found several accounts, the name is not correct etc.) you can write them down here.
7. If you did not find the Twitter account write down "Not found" in the column Twittername
8. If you encounter any names in the list that are not correct or if you noticed that a certain politician is not in the list, please let me know.





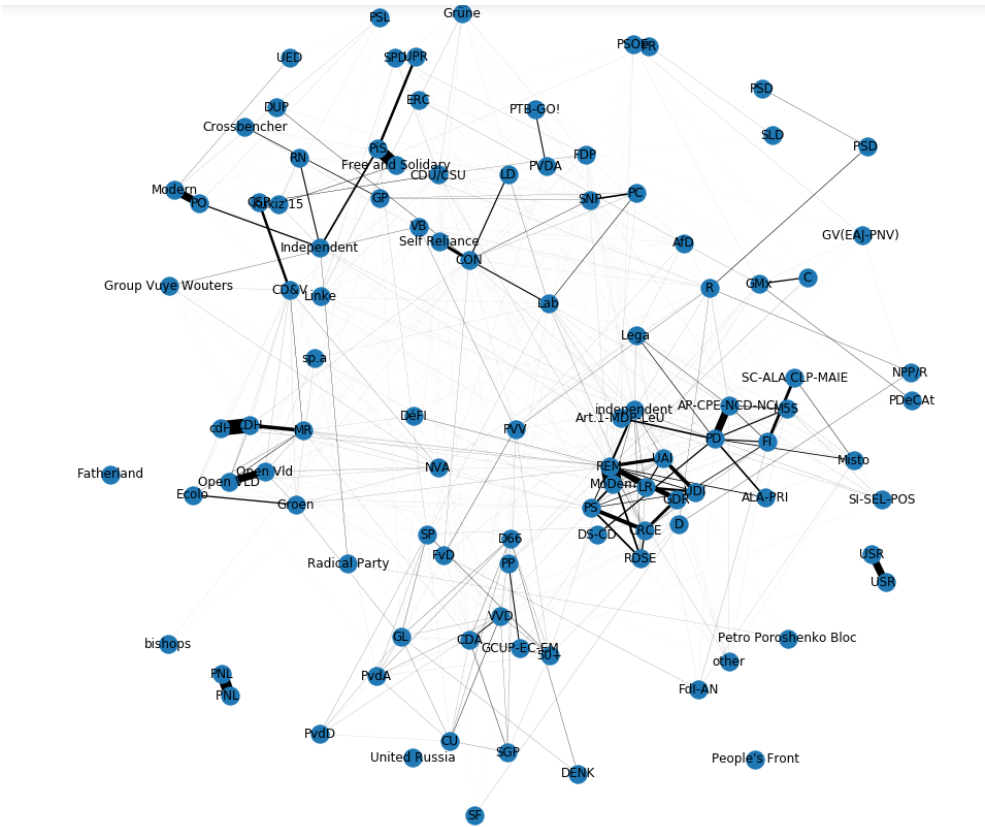


Figure 9.2: The parliamentary retweets network at party level.

9.2.3 Community detection

Table 9.7: Number of detected communities and adjusted mutual information between detected communities and parties.

Country	Parties	Followers			Mentions			Retweets		
		Modularity	Clusters	AMI	Modularity	Clusters	AMI	Modularity	Clusters	AMI
Italy	28	0.37	7	0.26	0.35	20	0.17	0.56	21	0.24
Belgium	18	0.31	5	0.59	0.28	5	0.65	0.66	8	0.78
Netherlands	14	0.13	6	0.33	0.20	7	0.44	0.61	15	0.71
Ukraine	13	0.19	8	0.07	0.47	8	0.10	0.52	7	0.07
UK	12	0.36	5	0.33	0.23	6	0.46	0.44	5	0.52
France	12	0.26	4	0.35	0.22	7	0.42	0.28	9	0.50
Romania	12	0.46	8	0.09	0.67	14	0.09	0.58	14	0.08
Spain	9	0.47	6	0.49	0.40	6	0.68	0.69	5	0.70
Germany	8	0.40	10	0.70	0.35	10	0.72	0.71	15	0.71
Poland	8	0.32	4	0.45	0.21	5	0.39	0.49	4	0.47
Russia	7	0.14	7	0.04	0.75	10	0.01	0.46	10	0.02
US	4	0.32	16	0.51	0.31	51	0.15	0.44	86	0.23

9.2.4 *Inter-party communication and ideology*

Table 9.8: RILE scores for the intersection of parties in our study and in the Manifesto Project Dataset (Volkens et al., 2020).

Country	Party	Rile	Country	Party	Rile
Belgium	PVDA	-33.681	Poland	PO	-13.31
	Groen	-21.849		PSL	-7.26
	sp.a	-19.199		K15	-1.77
	CD&V	-11.903		Modern	4.35
	OpenVLD	-8		PiS	10.81
	NVA	4.78	Romania	USR	-25.05
	VB	8.387		PSD	-17.81
France	FI	-30.019		UDMR	-15.28
	PS	-28.947		PMP	-1.01
	MoDem	-17.92		PNL	1.47
	PCF	-16.667		ALDE	26.67
	PRG	-10.056	Russia	SR	-25.19
	EÉLV	-8.636		KPRF	-18.30
	FN	1.674		ER	2.79
	UDI	13.619		LDPR	13.69
Germany	LR	13.619	Spain	ERC	-30.34
	LINKE	-41.914		PSOE	-29.27
	SPD	-21.437		PNV/EAJ	-11.63
	90/Greens	-21.058		C	-10.54
	FDP	0.578		PP	6.06
	CDU/CSU	2.757	UK	Labour	-31.85
Italy	AfD	17.43		SNP	-24.46
	D	-8.268		SF	-24.41
	M5S	-7.429		GPEW	-20.37
	Lega	4.656		LibDems	-19.61
	FDI	7.692		PC	-18.72
Netherlands	FI	15.625		Conservatives	6.21
	50PLUS	-31.11		DUP	12.26
	DENK	-24.83	Ukraine	Fatherland	-8.33
	SP	-23.04		Svoboda	0.00
	PvdD	-18.85	US	D	-20.58
	PvdA	-13.84		R	32.97
	GL	-9.35			
	D66	-6.54			
	CDA	3.60			
	CU	5.48			
	VVD	10.95			
	FvD	16.47			
	PVV	20.00			
	SGP	24.71			

### 9.3 Facebook likes to study lifestyle politics

#### 9.3.1 *Data collection*

The data collection took place in two waves in March and June 2018. In the first wave (March 2018), a detailed survey with questions on socio-demographics, media consumption, political preference and attitudes was sent to a representative panel of around 4,500 respondents. From these respondents, 524 agreed to give us access to their Facebook “like” data, via Facebook Login. In the second wave (May–June 2018), our shorter version of the survey was disseminated through the online webpages of popular Flemish newspapers to reach a broad audience. Another 6,209 participants agreed to give us access to their Facebook “likes”.

##### 9.3.1.1 *Survey invitation*

In wave 1, this survey invitation was sent to a representative panel (translated from Dutch to English):

“People increasingly keep themselves up to date on the news via social media such as Twitter and Facebook. Therefore, as researchers, we are very interested in the public pages you follow (like) on Facebook and the posts you share or post on your timeline. In addition, we want to examine what companies (such as Cambridge Analytica) can deduct from the like behavior of Facebook users and provide users insight into this.

For this we need your help. By participating in our study, you allow us access to your Facebook data and will need to answer ten questions. We will use those questions to determine what we can predict of people on the basis of their page likes on Facebook.

Of course, all data will be processed and stored fully anonymously. This research pursues scientific aims only, and data will not be shared with others. Your privacy is of high importance to us. Read our privacy policy here.

If you have any further questions about this research you can read our frequently asked questions section. If you still have any questions left you can contact us via [nwsdata@uantwerpen.be](mailto:nwsdata@uantwerpen.be).”

In wave 2, this text was added to the previous invitation<sup>1</sup>:

“Based on your Facebook likes we use data mining models to predict your gender, age, ideological leaning and party preference. We will also show you which specific likes (i.e. the Facebook pages that you have liked yourself) are the most important

<sup>1</sup> The original webpage can be found here (only in Dutch): <https://www.uantwerpen.be/nl/projecten/nws-data/facebookstudie/dataverzameling/>

contributors to this prediction. Do you want to find out what your Facebook likes reveal about you? You can test it here!"

#### 9.3.1.2 *Survey questions*

1. **How often did you use the following channels in the past month to follow news about current events? Never, less than once a week, once or twice a week, 3 or 4 times a week, daily, several times a day** (a) Radio (b) TV (c) Online (d) Newspapers (e) Facebook (f) Twitter
2. **To what extend are you interested in the following news subjects on a scale of 1 (not interested at all) to 5 (very strongly interested)?** (a) International (b) Politics (c) Local (d) Finance and economics (e) Entertainment (f) Lifestyle (cooking, fashion, travel) (g) Arts and culture (h) Sports (i) Science and technology (j) Justice and safety (k) Bizarre/funny
3. **In politics we often use the terms "left" and "right". Where would you situate your own opinion on a scale from 0 to 10, where 0 means "left", 5 "center", and 10 "right"?**
4. **Please indicate on a scale from 1 to 10 how likely it is that you will ever vote for each of the parties listed below. 1 means that you will never vote for this party and 10 means that you will definitely vote for this party.** (a) CD&V (b) Groen (c) N-VA (d) Open VLD (e) PVDA (f) Sp.a (g) Vlaams Belang
5. **To what extend are you interested in politics in general on a scale of 0 (not interested at all) to 10 (very much interested)?**
6. **What is your gender?** Male, female
7. **What is your age?** Younger than 25, between 25 and 55, older than 55
8. **What is your highest level of education?** No education, lower education, general education, technical education, vocational education, higher education (non-university), University education

Table 9.10: The number of participants in our dataset per target variable.

All participants (6.733)		
Leaning	Left	2,196 (33%)
	Center	2,322 (34%)
	Right	1,958 (29%)
	No answer	257 (4%)
Party preference	PVDA	750 (11%)
	Sp.a	781 (12%)
	Groen	1,988 (30%)
	CD&V	620 (9%)
	Open VLD	1,150 (17%)
	N-VA	1,270 (19%)
	VLaams Belang	150 (2%)
	No answer	24 (0%)

#### 9.3.1.3 Survey weights

We used Iterative proportional fitting (IPF)<sup>2</sup> to adjust survey weights to reflect the overall population distribution in terms of gender, age and education levels (see Table 9.11). However, the weighting process did not influence the outcomes all that much (e.g. the spearman rank correlation between the coefficients of the models with and without survey weights applied was 0.94) and therefore the unweighted results will be reported in the manuscript.

<sup>2</sup> We used the Python package `ipfn` <https://pypi.org/project/ipfn/>

Table 9.11: Gender, age and education of the Belgian population (StatBel, 2018)

	Variable	Survey total	Population percentage	Weighted survey total
Gender	male	3,747	50%	2,868
	female	2,078	50%	2,957
Age	younger than 25	1,598	28%	1,645
	between 25 and 55	3,558	40%	2,311
	older than 55	669	32%	1,868
Education	no education	75	3%	147
	lower education	41	5%	307
	vocational education	190	9%	550
	technical education	578	12%	679
	general education	937	16%	932
	higher education	1,661	18%	1,022
	university education	2,343	38%	2,188

### 9.3.2 Association methods

Association methods analyze the relationships between independent variables (e.g. Facebook pages) and a target variable (e.g. political leaning or party preference). In the association rule mining literature, several ‘interestingness measures’ are defined that discover relevant association rules of the form  $X \rightarrow Y$  (Jalali-Heravi and Zaïane, 2010; Geng and Hamilton, 2006; Kirchgessner et al., 2016). We will compare five bivariate methods: lift, leverage, binomial probability, chi-square and entropy; and one multivariate method: the coefficients of a logistic regression. We want to find the best (combination of) method(s) to rank Facebook pages according to their positive relationship with the target. For clarity, whenever we state that a certain Facebook page is left (or right), this means that the participants in our study who liked the Facebook page indicate themselves to be more left (or right) leaning. In the following sections,  $P$  denotes the probability and  $C$  the observed frequency or count.  $P(A, B)$  denotes the probability of  $A$  AND  $B$  both occurring. We will use the target variable *being left* as an example to describe the methods.

Table 9.12: Example of a contingency table for liking page X and being left.

	left	not left	TOTAL
like X	1	0	1
not like X	2.184	4.548	6.732
TOTAL	2.185	4.548	6.733

**LIFT** Lift is the degree to which two variables (often called antecedent and consequent or condition and result) are dependent on one another (McNicholas et al., 2008). It measures how much more prevalent the consequent is in the

selected sub population (based on the condition) compared to its prevalence in the total population. The lift of the rule *liking X implies being left* is defined as (Provost and Fawcett, 2013):

$$Lift(likeX \rightarrow left) = \frac{P(left|likeX)}{P(left)} = \frac{P(likeX, left)}{P(likeX)P(left)} \quad (9.1)$$

If liking X and being left are independent, the lift will be equal to 1. If the lift is greater than 1 the two variables are dependent on one another, and a rule based on the antecedent is potentially useful for predicting the consequent (Tufféry, 2011). To find the Facebook pages with highest association, the pages can be ranked from high to low lift. A downside of lift is that it does not take the frequency of antecedents into account and may find very strong associations for less frequent items. For example page X in Table 9.12 is liked only once, by a left user, and so the lift for the rule *liking X implies being left* receives the maximum value:

$$Lift(likeX \rightarrow left) = \frac{P(left|likeX)}{P(left)} = \frac{1/1}{2.185/6.733} = 3,08$$

In order to avoid infrequent pages to rank high, a minimum frequency can be set.

**LEVERAGE** Whereas lift is the ratio of two numbers, leverage is the difference between these numbers (Provost et al., 2015):

$$Leverage(likeX \rightarrow left) = P(likeX, left) - P(likeX)P(left) \quad (9.2)$$

If two variables are independent the leverage will be equal to 0. For our example, leverage is  $1/1 - 2.185/6.733 = 0,68$ . Unlike lift, leverage tends to prioritize antecedents with higher frequencies in the dataset. Facebook pages can again be ranked from high to low leverage.

**CHI-SQUARE STATISTIC** The dependence between liking X and being left can also be expressed based on the chi-square statistic. The statistic compares the observed values ( $O_i$ ) to the expected values in the contingency table if the variables were independent ( $E_i$ ):

$$\chi^2 = \sum \frac{[O_i - E_i]^2}{E_i} \quad (9.3)$$

With

$$E_i = \frac{[row\ total \times column\ total]^2}{total} \quad (9.4)$$

The Facebook pages with the highest chi-square values have the strongest association with a certain political preference. The statistic can be compared to

a chi-square distribution to determine the significance level of the rule (Jaiswal and Agarwal, 2012). The measure does not distinguish between a positive or negative relationship. For our example the expected values are given in Table 9.13 and the chi-square statistic is:

$$\chi^2 = \frac{[1-0,32]^2}{0,32} + \frac{[0-0,68]^2}{0,68} + \frac{[2184-2185,68]^2}{2185,68} + \frac{[4548-4.547,32]^2}{4.547,32} = 2,08$$

The p-value is the probability that a chi-square statistic with 1 degree of freedom is more extreme than 2.08, and is in this example equal to 0,15. In this case we cannot reject the hypothesis that the variables are independent.

Table 9.13: Expected values for the example of liking page X and being left.

	left	not left
like X	$1 \times 2.185/6.733 = 0,32$	0,68
not like X	2.185,68	4.547,32

**BINOMIAL PROBABILITY** The binomial distribution is used to obtain the probability of observing  $x$  successes in  $n$  independent experiments with the probability of success  $p_s$ . With the cumulative binomial distribution the probability of more than  $x = C(\text{like}X, \text{left})$  people liking a certain Facebook page out of the  $n = C(\text{left})$  people with a certain political preference and the independent probability  $p_s = P(\text{like}X)$  to like the Facebook page can be calculated as (Croarkin et al., 2002):

$$P(X > x) = 1 - \text{binomcdf}(x, n, p_s) \quad (9.5)$$

For our example, with  $x = 1$ ,  $n = 2.185$  and  $p_s = 1/6.733$ , the probability  $P(X > x)$  is 0,04. The Facebook pages with the smallest binomial probabilities have the strongest association with a certain political preference.

**ENTROPY** Entropy is a measure of disorder that captures how mixed (impure) a set is with respect to the properties of interest. For two classes *left* and *not left* with relative percentage  $p_1 = P(\text{left}|\text{like}X)$  and  $p_2 = P(\text{notleft}|\text{like}X)$  within the sub population of people who liked page X, entropy is defined as (Provost and Fawcett, 2013):

$$\text{entropy} = -p_1 \log(p_1) - p_2 \log(p_2) \quad (9.6)$$

For our example entropy is undefined because  $p_1 = 1/1$  and  $p_2 = 0$  and the logarithm of zero is undefined. Entropy ranges from close to zero at minimal disorder (the set contains almost exclusively members of the same class) to one at maximum disorder (the classes in the set are balanced with 50% class *left* and 50% class *not left*). Facebook pages can be ranked from low to high entropy. Again, entropy does not take the frequency of Facebook pages into account, so a minimum frequency must be set. Like the chi-square statistic, entropy also does not distinguish between a positive or negative relationship.



**COEFFICIENTS OF LOGISTIC REGRESSION MODEL** The logistic regression model is described in Chapter 2. Logistic regression models the ratio of chances  $P(left)$  and  $P(not\ left)$  of two possible outcomes, based on the independent variables  $x$  (i.e. the Facebook pages) (Sperandei, 2014):

$$\ln\left(\frac{P(left)}{1 - P(left)}\right) = \beta_0 + \beta_1 x_1 + \dots + \beta_k x_k \quad (9.7)$$

or:

$$P(left) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x_1 + \dots + \beta_k x_k)}} \quad (9.8)$$

The coefficients of the model indicate how much more (un)likely it is for the outcome to be present for instances with  $x = 1$  compared to the instances with  $x = 0$ . Therefore, the higher the coefficient of the variable in the model, the higher the positive association with the target. To avoid overfitting, an extra constraint is introduced to the optimization function that penalizes the weights of coefficients: regularization (Tibshirani et al., 2015). The most frequently-used regularization techniques are L1 regularization and L2 regularization.

While in traditional models the association strength of the included values can be directly derived from the ‘p-values’, the adaptive nature of the estimation procedure in regularized models makes this difficult. Assumptions about the asymptotic distribution of parameters when using a regularization term do not apply and therefore different methods for significance testing of the coefficient are suggested in literature (Lockhart et al., 2014; Tibshirani et al., 2015). We will follow the bootstrap procedure as described in Tibshirani et al. (2015). From the original dataset we take a random sample with replacement and built a model from this dataset to estimate the betas. This step is repeated 1000 times to obtain 1,000 values for each beta. For each beta we estimate the probability density function using a Gaussian kernel <sup>3</sup> to calculate the probability (p-value) that the parameter is less than or equal to zero. Facebook pages will be ranked based on the mean coefficients over 1,000 bootstraps while the p-value indicates significance on a 0.05 level.

To compare these six methods, we have distinguished three important characteristics for a method in order to provide insightful results: does the technique (a) take into account the number of instances that have a non-zero value for the feature (coverage, see Figure 9.3), (b) discriminate between a positive or negative association, and (c) score each top-ranked page with a unique value to produce a full ranking and reduce the number of ex aequos. Next, we will discuss the techniques in terms of these characteristics (see Table 9.14).

From Figure 9.3 we can conclude that both lift and entropy have a very low coverage. Indeed, lift does not consider the total number of instances that have liked a page.

<sup>3</sup> We used `gaussian.kde` from `scipy.stats` (Jones et al., 2001) with the default Scott’s Rule (Scott, 2015) for bandwidth selection.

If a page is liked by only one person and this happens to be someone with a left political leaning, this page will receive the highest lift score. Because of this, the infrequent Facebook pages often receive high lift scores and many pages receive the same score. This is also the case for entropy, which, moreover does not distinguish between a positive or negative relation. The ten most related pages when using entropy are all negatively related to being left: if you like one of these pages it is unlikely that you have a left political leaning. Leverage on the other hand seems to overvalue the coverage of the Facebook pages and therefore ranks very frequent pages too high. For example the page *Canvas* (TV broadcaster) only has a slight positive association with the target but is ranked in the top ten because of it is liked by many instances. The chi-square statistic (which can be used for significance testing) is based on the observed frequency versus the expected frequency and therefore considers the frequency of Facebook pages, but again the association can be positive or negative. The pages *Theo Francken* and *Bart De Wever* (both leading politicians of the center-right N-VA) for example are negatively related to a left political leaning. The binomial probability can in our setting be interpreted as the chance that liking the page is not positively related with a left political leaning and can also be used for significance testing. This measure takes both frequency and the sign of the association into account but a drawback is that often several pages receive an equal value (e.g. when ranked, the first 22 pages have the exact same value) and thus need to be sorted according to a second criterion. Finally, the coefficients of a logistic regression model assess the strength of association of the Facebook pages to the target variable in interaction with all other Facebook pages. The sign of the coefficient indicates whether the association with the target is positive or negative. The regularized versions of logistic regression also have a high coverage (coefficients of non-regularized logistic regression (NR) have a very low coverage).

Based on this empirical and theoretical evaluation, we suggest to rank the pages based on the coefficients of a regularized logistic regression. This technique has a high coverage, discriminates between a positive and negative association and provides unique values.

Table 9.14: Comparison of the methods and their characteristics.

	Coverage	Sign of association	Unique values
Lift	No	Yes	No
Leverage	Yes	Yes	Yes
Chi-square	Yes	No	Yes
Binomial	Yes	Yes	No
Entropy	No	No	No
Coefficient (L1)	Yes	Yes	Yes
Coefficient (L2)	Yes	Yes	Yes
Coefficient (NR)	No	Yes	Yes




Figure 9.3: The number of instances with minimum one of the top  $k$  features for different association methods.

### 9.3.3 Statistical comparison of models

We use a corrected paired differences  $t$ -test (based on 10-fold cross-validation results) to compare the AUC of the different models (Nadeau and Bengio, 2000) and apply the Bonferonni correction for multiple pairwise comparison (Vázquez et al., 2001; Pizarro et al., 2002). The results can be found in Table 9.15 and Table 9.16.

Table 9.15: Results for the corrected paired  $t$ -tests for ideological leaning.

		T score	P-value
M4	M1	5.54	0.000*
M4	M2	13.64	0.000*
M4	M3	31.64	0.000*
M1	M2	9.02	0.000*
M1	M3	21.42	0.000*
M2	M3	8.27	0.000*
Left	Right	4.76	0.000**
Left	Center	21.85	0.000**
Right	Center	14.19	0.000**

\* Significant with  $\alpha/6 = 0.008$

\*\* Significant with  $\alpha/3 = 0.017$

Table 9.16: Results for the corrected paired t-tests for party preference.

		T score	P value
M4	M1	-0.103	1.082
M4	M2	16.338	0.000*
M4	M3	40.747	0.000*
M1	M2	17.018	0.000*
M1	M3	22.860	0.000*
M2	M3	13.889	0.000*
N-VA	Vlaams Belang	1.754	0.079
N-VA	Open VLD	7.101	0.000**
N-VA	PVDA	8.488	0.000**
N-VA	Groen	11.771	0.000**
N-VA	Sp.a	8.401	0.000**
N-VA	CD&V	13.311	0.000**
Vlaams Belang	Open VLD	1.193	0.233
Vlaams Belang	PVDA	1.104	0.270
Vlaams Belang	Groen	1.890	0.059
Vlaams Belang	Sp.a	3.099	0.002
Vlaams Belang	CD&V	2.946	0.003
Open VLD	PVDA	0.051	0.960
Open VLD	Groen	1.953	0.051
Open VLD	Sp.a	2.639	0.008
Open VLD	CD&V	3.816	0.000**
PVDA	Groen	1.784	0.075
PVDA	Sp.a	2.481	0.013
PVDA	CD&V	3.156	0.002
Groen	Sp.a	1.969	0.049
Groen	CD&V	2.861	0.004
Sp.a	CD&V	0.778	0.437

\* Significant with  $\alpha/6 = 0.008$ \*\* Significant with  $\alpha/21 = 0.002$

## 9.3.4 Voter profiles

Table 9.17: The ten most related Facebook pages (with exclusion of political pages) to ideological leaning offer interesting insights in voters' interests.

	Feature	Description	Mean coefficient	P-value
Left	De Morgen	News & media website	0.32	0.000
	De Wereld Morgen	News & media website	0.31	0.000
	HART BOVEN HARD	Citizens' initiative for more solidarity	0.26	0.000
	Apache	News & media website	0.20	0.000
	Amnesty International	Nonprofit organization - human rights	0.17	0.000
	Ringland	Citizens' initiative for green mobility	0.15	0.000
	Dagen Zonder Vlees	Citizens' initiative to consume less meat	0.13	0.000
	MO*	News & media website	0.12	0.000
	Vrolijk Relativerende Liga ter Bestrijding van Azijnpis & Verzuring	Playful page with subtle criticism toward society	0.12	0.000
	Ish Ait Hamou	Belgian dancer, choreographer, television presenter and author of Moroccan descent	0.11	0.000
Center	KU Leuven	University	0.09	0.003
	NoodweerBenelux	Weather forecasts	0.08	0.002
	Tasty	Food videos and recipes	0.08	0.011
	She.be	News & media website for women	0.08	0.005
	Ben & Jerry's	Ice cream	0.08	0.001
	Milow	Belgian musician	0.07	0.000
	The Economist	News & media website	0.07	0.004
	Avatar	Fantasy /Sciencefictionfilm	0.07	0.000
	Politie Leuven	Police station	0.07	0.000
	Radio 2	Radio station	0.07	0.002
Right	SCEPTR	News & media website	0.17	0.000
	HLN.be	News & media website	0.17	0.000
	De Tijd	News & media website	0.16	0.000
	Onafhankelijk Verbond Der Vlaamsche Meme	Flemish nationalistic memes	0.13	0.000
	De Fiere Vlaamse Meme	Flemish nationalistic memes	0.12	0.000
	Schild & Vrienden	Flemish nationalistic youth movement	0.12	0.000
	Vlaamse Volksbeweging	Nonprofit organization - pro-Flemish	0.12	0.000
	Duvel	Beer	0.10	0.001
	Club Brugge K.V.	Sports team	0.10	0.002
	Dan Bilzerian	Internet celebrity, poker player and actor known for his lavish lifestyle	0.10	0.000

Table 9.18: The significant questions to ideological leaning for the survey (M<sub>3</sub>).

	Feature	Mean coefficient	P-value
Left	Q2: news subjects - arts & culture	0.40	0.000
	Q6: gender - female	0.33	0.000
	Q2: news subjects - politics	0.26	0.000
	Q7: age - between 25 and 55	0.21	0.000
	Q2: news subjects - international	0.19	0.000
	Q8: education - university	0.08	0.045
	Q5: interest in politics	0.03	0.031
Center	Q2: news subjects - finance & economy	0.11	0.001
	Q2: news subjects - international	0.09	0.006
	Q8: education - technical	0.09	0.046
	Q1: news channels - TV	0.09	0.000
	Q8: education - vocational	0.08	0.045
	Q2: news subjects - arts & culture	0.07	0.009
	Q2: news subjects - arts & culture	0.07	0.009
	Q2: news subjects - arts & culture	0.07	0.016
	Q1: news channels - Twitter	0.06	0.001
	Q1: news channels - radio	0.04	0.012
Right	Q2: news subjects - finance & economy	0.33	0.000
	Q2: news subjects - justice & safety	0.29	0.000
	Q5: interest in politics	0.07	0.000
	Q2: news subjects - bizarre/funny	0.06	0.036
	Q1: news channels - TV	0.04	0.035

Table 9.19: The five most related NMF components to left ideological leaning (based on the coefficients of a logistic regression model) and the Facebook pages they consist of. The first and second component mainly consist of pages related to the green party and the workers' party respectively. Component three is related to the city of Ghent and component four to Antwerp. Finally, the fifth component groups news-related pages.

Component 1	Component 2	Component 3
Natuurpunt	Peter Mertens	DOK
Groen	Raoul Hedebouw - Nederlands	Ghent in Motion
Greenpeace Belgium	HART BOVEN HARD	Stad Gent
Klimaatzaak	De Wereld Morgen	Kunstencentrum Vooruit
EVA	PVDA	De Hipste adresjes van Gent
Dagen Zonder Vlees	ManiFiesta	Gent
Oxfam-Wereldwinkels	Kim De Witte	De Gentenaar
MO*	Tom De Meester	Universiteit Gent
HART BOVEN HARD	Raoul Hedebouw	KERK GENT
Netwerk Bewust Verbruiken vzw	MO*	Gent Jazz Festival
11.11.11	Apache	Paard Van Troje
Wouter Deprez	Vluchtelingenwerk Vlaanderen	Gentblogt
Vluchtelingenwerk Vlaanderen	Solidair	Lichtfestival Gent
Velt vzw	Bleri Lleshi	GENT - Most beautiful city in Belgium
Kristof Calvo	EPO Uitgeverij	STAM - Stadsmuseum Gent
Loesje	PVDA Antwerpen	Eva Mouton
De Wereld Morgen	PTB	Gentse Feesten
Bond Beter Leefmilieu	Bernie Sanders	Gentse feesten
Lekker bio	Jacobin Magazine	CirQ
Radio 1	Lava Tijdschrift	De Krook
Component 4	Component 5	
Stad Antwerpen	De Standaard	
Zomer van Antwerpen	VRT NWS	
De Roma	De Morgen	
Ringland	Canvas	
OLT Rivierenhof	Knack	
Bar Noord Zomerbar	Radio 1	
Plein Publiek	De Ideale Wereld	
Zomerfabriek	Studio Brussel	
De Studio	In het spoor van Rudi Vranckx	
ZOO Antwerpen	De Tijd	
Trix	Sporza	
UAntwerpen	Humo	
deSingel Internationale Kunstcampus	De Wereld Morgen	
Fotomuseum Antwerpen - FOMU	Groen	
Toneelhuis	Eén	
Plantenasiel Antwerpen	Terzake	
Mercado	Apache	
MAS — Museum aan de Stroom	Nieuwsblad.be	
EcoHuis Antwerpen	MO*	
Gazet van Antwerpen	De Afspraak	

Table 9.20: The five most related NMF components right ideological leaning. The first and second component group pages related to the Flemish nationalists party and the liberal party respectively. The third component groups local branches of the Flemish nationalists party. The fourth component consists of soccer players and finally, the fifth component contains (electronic dance) music.

Component 1	Component 2	Component 3
Bart De Wever	Open Vld	Jong N-VA Beveren
Theo Francken	Guy Verhofstadt	Jong N-VA Hamme
N-VA	Alexander De Croo	Jong N-VA Kortenberg
Tomas Roggeman	Gwendolyn Rutten	Jong N-VA Wetteren
Jong N-VA	Jong VLD	N-VA Hoeilaart
Ben Weyts	Maurits Vande Reyde	Jong N-VA Geel
Vlaamse Volksbeweging	ALDE Party Liberals and Democrats for Europe	Jong N-VA Beringen
Lawrence Vancraeynest	European Parliament	Jong N-VA Brecht
Lorin Parys	ALDE Group	Jong N-VA Zuiderkempen
Zuhal Demir	Maggie De Block	Jong N-VA Pajottenland
Geert Bourgeois	Philippe De Backer	Jong N-VA Haacht
SCEPTR	Bart Tommelein	Jong N-VA Geraardsbergen
Doorbraak	Charles MICHEL	N-VA Meise-Wolvertem
Sander Loones	Mathias De Clercq	Jong N-VA Hoeselt
Peter Dedecker	Hans Maes	JONG N-VA Zemst
Christiaan Janssens	Annemie Turtelboom	Jong N-VA Beernem
Nabilla Ait Daoud	De Tijd	Jong N-VA Lokeren
Eva Paelinck	VRT NWS Politiek	Jong N-VA Ekeren
Piet De Bruyn	Patrick Dewael	Jong N-VA Zuidrand
Yoleen Van Camp	Bart Somers	Jong N-VA Noord-Limburg
Component 4	Component 5	
Belgian Red Devils	Korsakoff	
Vincent Kompany	Da Tweekaz	
Sporza	Mark With a K	
Thibaut Courtois	Wildstylez	
Kevin De Bruyne	Brennan Heart	
Eden Hazard	Psyko Punkz	
Romelu Lukaku	The Qontinent	
Dries Mertens	Ran-D	
Play Sports	Zatox	
Marouane Fellaini	Bass Events	
Jupiler Pro League	Angerfist	
Sportwereld	Headhunterz	
Belgian Football	Noisecontrollers	
Divock Origi	Dirty Workz	
Thomas Vermaelen	D-Block & S-te-Fan	
VoetbalNieuws.be	Q-dance	
Axel Witsel	The Oh!	
FIFA World Cup	Frontliner	
Leo Messi	Evil Activities	
FC Barcelona	Gunz for Hire	



## 9.4 Polarization in lifestyle domains

### 9.4.1 Sample details

Table 9.21: Comparison of samples.

	Full Sample	Sample 2 <sup>a</sup>	Sample 3 <sup>b</sup>	<i>p</i> <sup>c</sup>	Sample 4 <sup>d</sup>
% Democrat	31	32	40	0.17	46
Mean ideology (5-point)	2.98	2.89	2.76	0.01	2.79
Mean news interest (1-4)	3.28	3.27	3.32	0.10	3.33
Mean age	51	49	49	0.16	48
% High school or less	23	20	22	0.17	26
% Female	54	57	57	0.67	58
% Less than \$20,000	14	14	14	0.82	15
% White	75	77	76	0.49	76
Median number of page likes					374
Median number of page likes with minimum 30 likes					115
<i>N</i>	3,500	2,711	1,331		1,211

<sup>a</sup>Respondents who said in the survey that they have a Facebook account (i.e., they selected “Facebook” from the list of response options to the question “Do you have accounts on any of the following social media sites?”).

<sup>b</sup>Respondents (regardless of their answer in the previous question) who consented to share Facebook profile information with the researchers.

<sup>c</sup>*p*-values are computed from *t*-tests of the difference in means between the sample of respondents who reported having a Facebook account and those who consented to provide access to their profile data.

<sup>d</sup>The final column subsets to those who shared any Facebook data at all that we were able to link back to the survey and that have liked at least one of the pages included in our analysis.

Table 9.22: Total number of Page Likes (for the 5155 pages included in our study) per ideology

Ideology score	Number of likes	Ideology	Number of likes
Very liberal	47,440 (17%)	Liberal	105,473 (38%)
Liberal	58,033 (21%)		
Moderate	96,837 (34%)	Moderate	96,837 (34%)
Conservative	41,384 (15%)	Conservative	78,466 (28%)
Very conservative	37,082 (13%)		

#### 9.4.2 Coding categories

Starting from the initial Facebook categories<sup>4</sup>, we extracted a list of 24 categories. Two independent coders received the following instructions and a short training session to assign Facebook pages to the corresponding category. After 500 and 1000 coded pages we evaluated inter coder agreement and refined the instructions where needed.

**INSTRUCTIONS** The goal of this coding task is to divide Facebook pages into 24 categories, based on the content of the Facebook page and your background knowledge about the subject. The 24 categories can be found in the Codebook (Table 9.23), together with a description and some illustrative examples. Some things to keep in mind while coding:

- Select the correct category from the drop-down list. You can assign multiple categories to one Facebook page, but try to indicate the primary category first. For example, The Daily Show is a TV Show in the first place, but could also be classified as 'Entertainment' or even 'News & Media'.
- Make sure to visit the Facebook page (by following the link provided) before you assign a category. For example, Harry Potter could refer to both the books or the movies and you will need to visit the page in order to assign the correct category.
- If the link we provided is broken, please try to find the page manually. If that does not work, the page is likely removed from Facebook. We ask that you indicate this in the 'Error' column and that you try to assign a category based on the name of the page alone.
- If a page does not belong to any of the provided categories you can assign the category 'Other'. In this case, we ask you to specify in your own words which category you would assign to this page.

<sup>4</sup> see <https://www.facebook.com/pages/category/>

After coding, the smallest categories were grouped together to result in the final categorization that was shown in Table 6.2. **Religion** and **Identity/affinity groups** were merged to **Identity & religion**, **Books** were added to **Arts & Culture** and **Bars & Nightlife** to **Entertainment**. Thereafter, inter coder reliability was calculated. For the final results, the codings of one of the two coders were selected (after a shallow quality inspection), to ensure internal consistency of the coded categories. Finally, pages in the category **Shows & Events** were manually reassigned by the authors to a corresponding category (e.g. Music, TV Shows, Movies or Entertainment) and also pages from the category **Other** were assigned to a suitable category. (Non-) scientific research was added to **Education** and the category was renamed to **Research & Education**.

Table 9.23: Categories codebook

Category	Description	Examples
Government & Politics	Politicians, political parties, political content, political communities and government organizations	Barack Obama, Being Conservative, U.S. Army
News & Media	News, media, radio, magazines, journalists, etc.	Fox News , The Economist
Individual opinion leaders	Individual influencers, bloggers, commentators, etc.	Michelle Malkin, Michael Moore
Civil Society	Nonprofit organizations and labor unions (formal organizations)	Human rights campaign, AFL-CIO
Religion	Religious pages, religious organizations	Jezus loves you, Franklin Graham,
Identity/affinity groups	Pages referring to home country, region, ethnic or cultural groups	Israel is my heart, Africans-In-America
Books	Books, libraries, publishers, writers, poetry, thematical magazines	Barnes & Noble, Lord of the Rings books, Stephen King
Tv Shows	TV shows, episodes, channels	The Big Bang Theory, The Daily Show, National geographic channel
Music	Music, bands, producers, record labels, albums, etc.	The Beatles, Gibson, Warner Music Group
Movies	Movies, actors, directors, movie characters, cinema and	Harry Potter (movie), Regal cinemas, Alfred Hitchcoc
Food & Beverage	Food, cooking, restaurants, drinks, spirits, breweries etc.	Starbucks, Pepsi, Tasty
Sports	Sports, teams, athletes, leagues, games, gym	New York Yankees, NFL, Road Runner sports
Beauty & Health	Cosmetics, healthcare, medical	MinuteClinic, NIVEA, Bayer Aspirin
Arts & Culture	Arts, culture, photography, museums, artists, musicals, theater, etc.	American museum of natural history, Andy Warhol, WICKED the musical
Education	Schools, universities, student organizations and education	LeapFrog USA, New York University, VINCI Schools
Travel	Travel, tour agencies and tourism	Southwest Airlines, Hilton Hotels & Resorts, Love GREAT Britain
Bars & Nightlife	Bars, cafes, pubs, clubs etc.	House of Yes, Smalls Jazz Club (also music), The Wayland
Shows & events	One-off or limited occurrences, such as festivals, shows, performances and concerts	Honda Stage, Ultra Music Festival, TomorrowWorld
Entertainment	Entertainment, games, humor, amusement, comedy etc.	Larry the cable guy, Grumpy cat memes, Candy Crush
Public Figures	Public figures	Ellen DeGeneres, Michelle Obama, Dave Ramsey
Interests	Interests, communities (informal) and hobbies	Dogs, Hippie Peace Freaks, Humans of New York
Shopping & retail	Apparel, accessories, clothing, fashion, consumer electronics, home decaration, stores, shopping mall, wholesale, etc.	Converse, Amazon.com, Bose
Cars and transportation	Car brands, automotive, airlines, boats, etc.	Hyundai, American Airlines, GasBuddy
Services	Marketing, advertising, legal, finance, consulting, etc.	BFAds - Black Friday Ads, PayPal, Facebook Business
Other	Everything that does not fit in the other categories. Please specify in your own words which category you would assign to this page	

### 9.4.3 *Methods and measures*

**MODULARITY** is the relative density of edges inside communities with respect to edges outside of communities (Newman, 2010):

$$Q = \frac{1}{2m} \sum_{ij} (A_{ij} - \frac{k_i k_j}{2m}) \delta(c_i, c_j), \quad (9.9)$$

where  $m$  is the number of edges,  $A_{ij}$  represents the weight of the edge between node  $i$  and  $j$ ,  $k_i$  is the number of edges adjacent to node  $i$ , and  $\delta(c_i, c_j)$  is 1 if  $i$  and  $j$  are in the same community and 0 otherwise. A modularity of zero indicates that the fraction of within-community edges is no different from what we would expect for a randomized network, while a value above about 0.3 is found to be a good indicator of significant community structure in a network (Clauset et al., 2004). By modularity optimization networks can be divided into communities or clusters. For example, the modularity of a network of 105 books on American politics on Amazon.com is  $Q = 0.52$  (Newman, 2006). In this network four communities were identified: one consisting almost entirely of liberal books and one almost entirely of conservative books, and two containing most of the centrist books. Similarly, a network of political 2004 U.S election blogs (Adamic and Glance, 2005) was cleanly divided into a conservative and a liberal community with an optimal modularity of  $Q = 0.43$  (Newman, 2006).

**CRAMER'S V** is a normalized version of the Chi-square statistic and determines the effect size. It is defined by (Cohen, 2013):

$$V(Z) = \sqrt{\frac{\chi(Z)/N^1}{\min(r-1, c-1)}} \quad (9.10)$$

where  $r$  is the number of rows (2) and  $c$  the number of columns (3) in the contingency table.

### 9.4.4 *Robustness checks*

#### 9.4.4.1 *Including all Facebook pages*

We only include pages that are liked by a minimum of 30 respondents in our analysis, to ensure that the calculated page ideology and homogeneity scores are reliable<sup>5</sup>. A very limited amount of pages –5155– has a minimum of 30 Likes, which implies that almost 99% of the Facebook pages in our dataset can not be used for analysis. Per user, on average only one third of their likes consists of pages with a minimum of 30 Likes. Although this seems like a huge loss of information, in Figure 9.4 and Figure 9.5 we show that the distribution of individual page ideology and

<sup>5</sup> With a sample size of 30 Likes, we can estimate page ideology with 95% confidence and a margin of error no larger than 0.4.

homogeneity scores based on all pages and based on only pages with a minimum of 30 Likes is in fact highly similar. The mean individual homogeneity score for all pages (0.097) is only slightly lower than for the pages with 30 Likes (0.105). Yet, we do acknowledge that the results and conclusions in this paper our based on widespread Facebook pages and that we cannot analyze polarization in smaller pages.

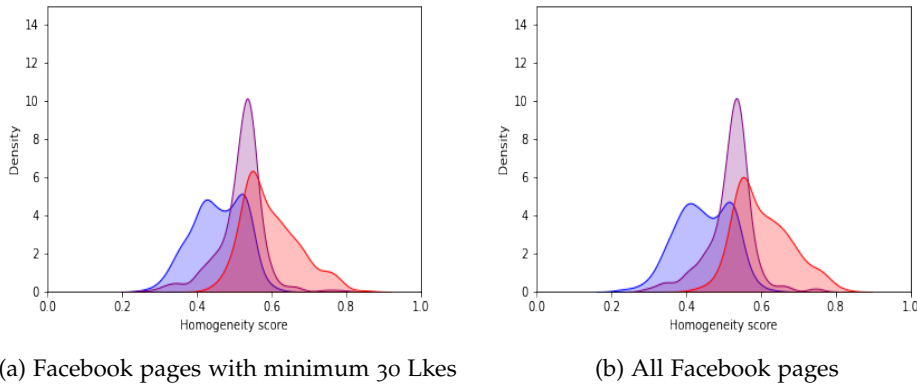


Figure 9.4: Page ideology distribution for liberals (blue), moderates (purple) and conservatives (red) when taking into account (a) Facebook pages with a minimum of 30 Likes and (b) all Facebook pages in our dataset.

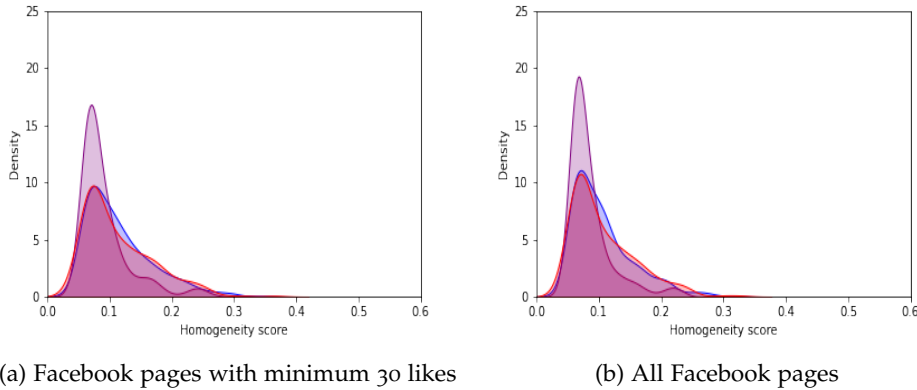


Figure 9.5: Homogeneity distribution for liberals (blue), moderates (purple) and conservatives (red) when taking into account (a) Facebook pages with a minimum of 30 Likes and (b) all Facebook pages in our dataset.

## 9.4.4.2 Homogeneity metrics

We acknowledge that several other metrics could be used to measure homogeneity (see Yamaya et al., 2020). As a robustness check, we compare our results using Cramer's V to two other measures of homogeneity: Entropy and Variance.

**Entropy** is a measure of disorder that captures how mixed (impure) a set is with respect to the properties of interest. Originally introduced by Shannon in 1994 to quantify uncertainty in strings of text (Shannon, 1948), it has become a basic quantity in information theory and has found its way into many applications. It is used as a basis for machine learning methods (Provost and Fawcett, 2013), for studying neuronal activity (Panzeri and Treves, 1996) or even as a measurement of biological diversity (Magurran, 2013). The Shannon entropy of Facebook page  $Z$  is defined as (Shannon, 1948):

$$S(Z) = - \sum_{c=0}^3 p(c) \log_2 c(k) \quad (9.11)$$

The probability  $p(c)$  that a user that liked page  $Z$  has ideology  $c$  is estimated using maximum likelihood, or  $p(c) = N_c^1 / N$ .

With three classes, entropy ranges from close to zero at minimal disorder (the page is liked by almost exclusively members of the same class) to 1.58 at maximum disorder (the classes are balanced with 33% class *liberal*, 33% class *moderate* and 33% class *conservative*). In other words, the lower the entropy, the more homogeneous the respondents that liked the page are in terms of their self-reported ideology. We will re-scale this value between 0 and 1.

**Variance** can also be used to measure audience homogeneity:

$$\sigma(Z) = \sum \frac{(s_i - I(Z))^2}{N} \quad (9.12)$$

with  $s_i$  the ideology score (0-4) of the  $i$ -th individual and  $I(Z)$  the average ideology of page  $Z$ .

The ranking of categories (see Table 9.24) is very similar for the three different metrics<sup>6</sup>. Apart from some slight individual differences, the overall conclusion stays the same. For the beta regressions (Table 9.25), the results for the dependent variables Entropy and Cramer's V are highly similar<sup>7</sup> but differ slightly from the regression using Variance.

<sup>6</sup> For Entropy and Variance, smaller values indicate higher homogeneity.

<sup>7</sup> Note that for Entropy and Variance the relation with homogeneity is negative, while for Cramer's V this is positive, hence the opposite sign for Cramer's V

Table 9.24: Entropy, Cramer's V, and Variance for the Facebook categories.

Category	Entropy	Cramer's V	Variance
Politics	0.66	0.22	0.05
Political news	0.72	0.18	0.06
Hardnews	0.79	0.15	0.07
Civil Society	0.80	0.12	0.08
Identity & Religion	0.86	0.12	0.09
Individual opinion leaders	0.81	0.12	0.08
Public Figures	0.84	0.10	0.08
Arts & Culture	0.87	0.07	0.09
Tv Shows	0.89	0.07	0.08
Entertainment	0.91	0.07	0.09
Research & Education	0.90	0.06	0.09
Music	0.91	0.06	0.09
Interests	0.91	0.06	0.09
Movies	0.90	0.06	0.09
Sports	0.91	0.05	0.08
Services	0.95	0.05	0.09
Beauty & Health	0.94	0.04	0.08
Travel	0.95	0.04	0.09
Shopping & retail	0.95	0.04	0.09
Food & Beverage	0.95	0.04	0.09
Cars and transportation	0.96	0.04	0.09
TOTAL	0.90	0.07	0.08



Table 9.25: Determinants of individual homogeneity for different homogeneity metrics.

	Cramer's V	Entropy	Variance
Age: 30-44	-0.002 (0.048)	0.073 (0.070)	-0.033** (0.013)
Age: 45-65	0.075* (0.045)	0.051 (0.067)	-0.064*** (0.013)
Age: Over 65	0.340*** (0.051)	-0.324*** (0.076)	-0.130*** (0.015)
Black	-0.144*** (0.044)	0.149** (0.066)	0.003 (0.012)
Hispanic	0.034 (0.052)	-0.048 (0.080)	-0.016 (0.016)
Other Race	-0.140*** (0.052)	0.184** (0.077)	0.028** (0.014)
Female	-0.117*** (0.026)	0.134*** (0.040)	0.035*** (0.008)
Income	0.001 (0.003)	-0.001 (0.005)	-0.001 (0.001)
Education	0.020** (0.009)	-0.058*** (0.014)	-0.003 (0.003)
Very Liberal	0.181*** (0.039)	-0.415*** (0.058)	-0.015 (0.012)
Liberal	0.086** (0.037)	-0.269*** (0.056)	-0.018* (0.011)
Conservative	0.172*** (0.036)	-0.115** (0.056)	-0.009 (0.010)
Very Conservative	0.231*** (0.046)	-0.160** (0.073)	0.003 (0.014)
Political news interest	0.158*** (0.018)	-0.249*** (0.027)	-0.035*** (0.005)
Number of likes	-0.00003*** (0.00001)	0.00001 (0.00001)	0.00001*** (0.00000)
Constant	-2.281*** (0.071)	1.955*** (0.107)	-2.430*** (0.020)
N	1,087	1,087	1,087
Pseudo R <sup>2</sup>	0.300	0.278	0.200

\*p &lt; .1; \*\*p &lt; .05; \*\*\*p &lt; .01

Beta regressions with survey weights applied. Reference categories are Age: 18-29, White race, Male gender, and Moderate ideology.

Income ranges from 1 to 31, Education from 1 to 6, and political news interest from 1 to 4.

#### 9.4.4.3 *Regressions*

As a robustness check, we include the regression results weighted for the number of likes in Table 9.26. The results are congruent with the regression results for Cramer's V in Table 9.25. Secondly we include party identification strength in stead of ideology in Table 9.27. We do not find a significant relation for party identification strength and liking more homogeneous pages for any of the groups. The explanation could be that users with moderate ideology are less likely than both liberals and conservatives to like homogeneous pages and since moderates more often identify as democrats than as republicans we find republicans to be more likely to like homogeneous pages, independent of party identification strength.

Table 9.26: Determinants of individual homogeneity when regression is weighted for individual number of page likes

	Homogeneity
Age: 30-44	-0.080*** (0.022)
Age: 45-65	0.068*** (0.021)
Age: Over 65	0.191*** (0.024)
Black	-0.102*** (0.019)
Hispanic	0.031 (0.023)
Other Race	-0.064*** (0.017)
Female	-0.209*** (0.012)
Income	0.001 (0.001)
Education	0.034*** (0.004)
Very Liberal	0.092*** (0.017)
Liberal	0.035** (0.015)
Conservative	0.120*** (0.016)
Very Conservative	0.323*** (0.018)
Political news interest	0.113*** (0.007)
Number of likes	-0.00001*** (0.00000)
Constant	-2.411*** (0.032)
N	1,087
Pseudo R <sup>2</sup>	0.270

\*p &lt; .1; \*\*p &lt; .05; \*\*\*p &lt; .01

Beta regression weighted for individual page likes. Reference categories are Age: 18-29, White race, Male gender, and Moderate ideology. Income ranges from 1 to 31, Education from 1 to 6, and political news interest from 1 to 4.

Table 9.27: Determinants of individual homogeneity with party identification (PID) strength

	Politics	Political news	Hardnews	Lifestyle
Age: 30-44	-0.110* (0.057)	0.185** (0.084)	-0.021 (0.093)	-0.040 (0.045)
Age: 45-65	-0.084 (0.054)	0.206*** (0.080)	0.029 (0.087)	0.005 (0.042)
Age: Over 65	-0.034 (0.060)	0.333*** (0.087)	0.223** (0.095)	0.232*** (0.048)
Black	0.032 (0.054)	-0.255*** (0.073)	-0.265*** (0.084)	-0.136*** (0.042)
Hispanic	-0.073 (0.066)	-0.106 (0.089)	-0.194** (0.099)	0.061 (0.050)
Other Race	-0.108* (0.064)	-0.313*** (0.087)	-0.382*** (0.098)	-0.093* (0.049)
Female	0.007 (0.032)	-0.040 (0.041)	-0.092** (0.046)	-0.120*** (0.025)
Income	0.007* (0.004)	-0.012** (0.005)	-0.002 (0.006)	0.001 (0.003)
Education	-0.012 (0.011)	0.024 (0.015)	0.035** (0.017)	0.023*** (0.009)
PID strength	0.028 (0.021)	0.036 (0.026)	0.037 (0.030)	0.014 (0.016)
Democrat	0.088 (0.070)	-0.128 (0.090)	-0.172* (0.102)	-0.020 (0.053)
Republican	0.295*** (0.067)	0.324*** (0.086)	0.189* (0.098)	0.073 (0.051)
Political news interest	0.084*** (0.021)	0.198*** (0.029)	0.232*** (0.033)	0.143*** (0.016)
Number of likes	-0.00002*** (0.00001)	-0.00001 (0.00001)	-0.00001 (0.00001)	-0.00003*** (0.00001)
Constant	-1.219*** (0.091)	-1.629*** (0.126)	-1.556*** (0.140)	-2.309*** (0.071)
N	847	793	760	1,117
Pseudo R <sup>2</sup>	0.100	0.210	0.204	0.250

\*p &lt; .1; \*\*p &lt; .05; \*\*\*p &lt; .01

Beta regressions with survey weights applied. Reference categories are Age: 18-29, White race, Male gender, and Independent.

Income ranges from 1 to 31, Education from 1 to 6, PID strength from 1 to 4, and political news interest from 1 to 4.

9.4.5 *Results*9.4.5.1 *Network analysis*

Network analysis of the bottom node projection points in the direction of political polarization on Facebook (see Figure 9.6). This is most outspoken when we consider only political pages and the community detection algorithm reveals two communities with a distinctly different ideological composition (Table 9.28). For political news pages and hard news a third (and even a fourth in case of the latter) community exists that is predominantly composed of moderate voters. For lifestyle pages however, the three detected communities show hardly any ideological differences, and the modularity of these communities is much lower than for the political pages (Table 9.28).

Table 9.28: Size, density and hierarchy of the communities in the bottom node projection.

Group	Modularity	Community	Ideology	Size	Density	Hierarchy
Politics	0.38	1	0.28	456	0.65	0.34
		2	0.7	344	0.52	0.45
Political news	0.30	1	0.29	389	0.45	0.53
		2	0.71	271	0.63	0.37
		3	0.52	86	0.45	0.39
Hardnews	0.25	1	0.28	306	0.56	0.40
		2	0.71	223	0.60	0.39
		3	0.49	138	0.39	0.60
		4	0.46	46	0.69	0.27
Lifestyle	0.15	1	0.46	786	0.63	0.34
		2	0.54	286	0.87	0.13
		3	0.40	5	0.70	0.30
All	0.16	1	0.34	555	0.74	0.26
		2	0.53	288	0.91	0.09
		3	0.73	241	0.81	0.19

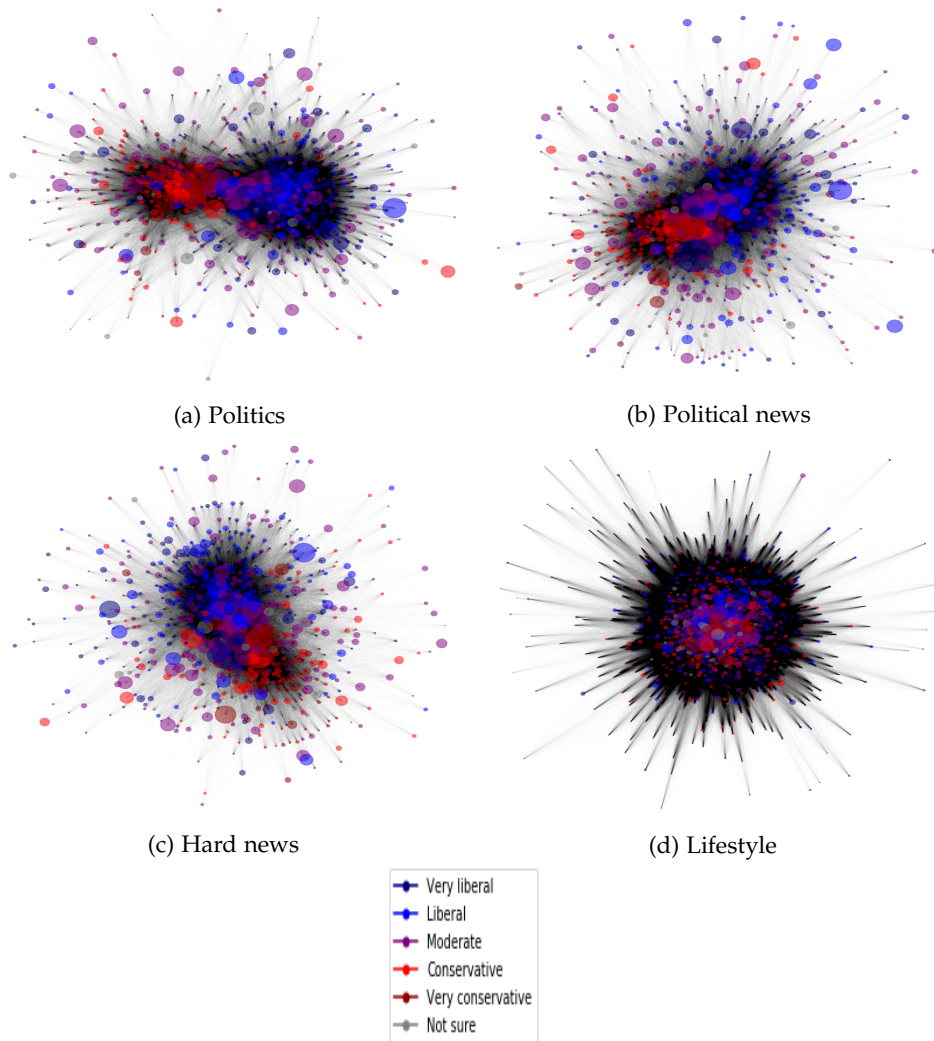


Figure 9.6: Network visualization for the bottom node projection of (a) Political pages, (b) Political news, (c) Hard news, and (d) Lifestyle pages. The size of the bubble represents the total number of likes of the user, the color represents the ideology of the user.

Table 9.29: Mean ideology, size, density and hierarchy of the communities in the top node projection.

Group	Modularity	Community	Ideology	Size	Density	Hierarchy
Politics	0.40	1	0.31	102	0.98	0.02
		2	0.80	131	0.98	0.02
Political news	0.28	1	0.80	72	0.98	0.02
		2	0.38	87	1.00	0.00
Hardnews	0.17	1	0.40	49	1.00	0.00
		2	0.56	9	1.00	0.00
		3	0.77	30	0.93	0.07
Lifestyle	0.13	1	0.50	1912	0.90	0.10
		2	0.50	1734	0.99	0.01
All	0.15	1	0.51	2429	0.89	0.11
		2	0.50	1986	0.99	0.01

9.4.5.2 Facebook categories

In Figure 9.7, we map the ideology score per Facebook page on the x-axis, and the homogeneity score on the y-axis. Note that page ideology and page homogeneity are related analytically. Pages with a very homogeneous liberal or conservative audience will accordingly have a very low or high page ideology score, while pages with a diverse audience will have a moderate page ideology score, hence the U-shaped graphs. A moderate page ideology score combined with high homogeneity score indicates a predominantly moderate audience (e.g. Serena Williams in Figure 9.7e), but these pages are rare. For homogeneous categories (e.g. Figure 9.7a) the majority of pages are located at both ends of the U-shaped graph, while for more heterogeneous categories (e.g. Figure 9.7f) the pages are located in the bottom-center of the graph.

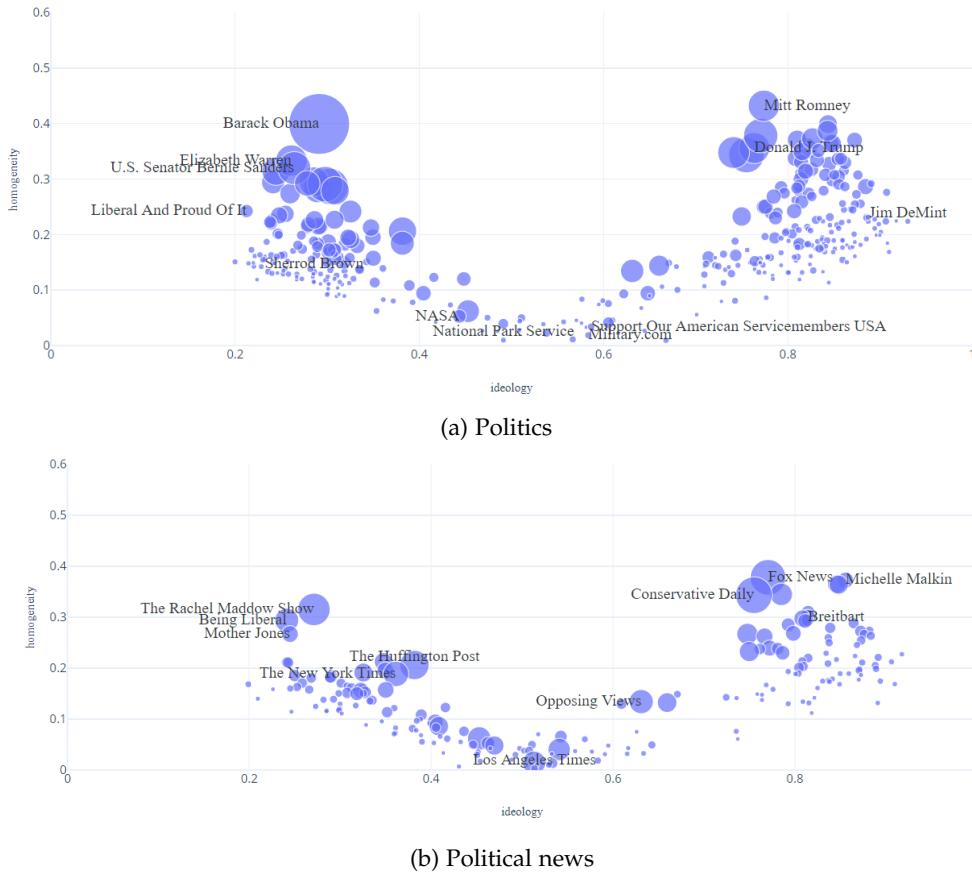




Figure 9.7: Ideology and homogeneity scores for the Facebook pages per category. The magnitude of the circle represents the total number of likes of the Facebook page.






(c) Hard news




(d) Arts &amp; Culture

Figure 9.7: (Continued) Ideology and homogeneity scores for the Facebook pages per category. The magnitude of the circle represents the total number of likes of the Facebook page.



(e) Sports



(f) Food &amp; Beverages

Figure 9.7: (Continued) Ideology and homogeneity scores for the Facebook pages per category. The magnitude of the circle represents the total number of likes of the Facebook page.

## 9.4.5.3 Facebook users

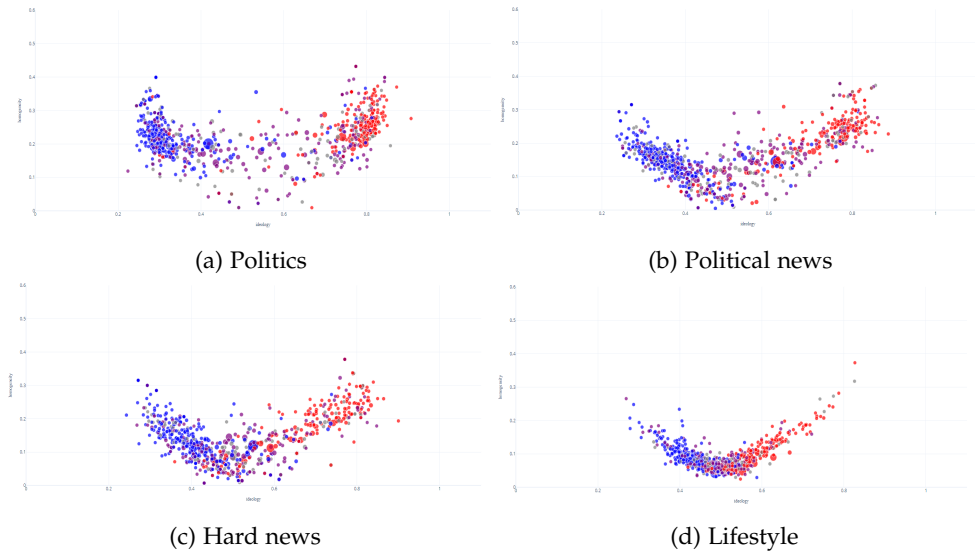


Figure 9.8: Page ideology and homogeneity of the users when taking into account (a) Political pages, (b) Political news, (c) Hard news, and (d) Lifestyle pages. The size of the bubble represents the total number of likes of the user, the color represents the self-reported ideology of the user.

Table 9.30: Overlapping coefficient (bootstrap mean) for the probability density curves of liberals (Lib), moderates (Mod) and conservatives (Con).

	Lib-Con	Lib-Mod	Con-Mod
Politics	0.15	0.64	0.47
Political news	0.24	0.69	0.48
Hard news	0.36	0.75	0.54
Lifestyle	0.75	0.89	0.85

## 9.5 Polarization in Belgium compared to the U.S.

Table 9.31: Overlapping coefficient (bootstrap mean) for the probability density curves of left, center, and right leaning respondents.

	Left-Right	Left-Center	Right-Center
Politics	0.36	0.62	0.59
News	0.72	0.86	0.85
Lifestyle	0.68	0.81	0.86

## Colophon

This document was typeset in  $\text{\LaTeX}$  using the typographical look-and-feel `classicthesis`. Most of the graphics in this thesis are generated using `pgfplots` and `pgf/tikz`. The bibliography is typeset using `natbib`.