The Measurement and Representation of Influencer Communities in United States Political Discourse on Social Media

Anna Beers

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2024

Reading Committee:

Kate Starbird, Chair

Emma Spiro, Chair

Benjamin Mako Hill

Program Authorized to Offer Degree:

Human Centered Design and Engineering

University of Washington

Abstract

The Measurement and Representation of Influencer Communities in United States Political Discourse on Social Media

Anna Beers

Chairs of the Supervisory Committee:

Kate Starbird
Department of Human Centered Design and Engineering

Emma Spiro
Information School

The United States (U.S.) is in the midst of a paradigm shift in who creates news. It is widely known that the internet has in the past few decades displaced television and print newspapers as the primary medium where news is consumed. However, the past two decades have also seen a shift in *who* is communicating the news. People consuming digital news, and especially young people, are now less likely than they were to get their news from journalists or television broadcasters, and more likely to receive it from a nebulous figure commonly referred to as the influencer. People are not only paying more attention to influencers, but there has seemingly been a remarkable growth in the type of and number of influencers in the last decade. This growth has been enabled both by the maturation of new, massively-networked social media platforms and the industrialization of influencing as a profession and an economic infrastructure.

In this dissertation, I use quantitative and qualitative methods to analyze case studies of U.S. political discourse networks to better understand the structure, tactics, and consequences of our new era of influencer media. After conducting a literature review of previous work on influencers in political communication and media studies, I propose flexible ways to represent both influencers and what I call influencer communities in social media networks. Influencer communities are assemblages of influencers working cooperatively and antagonistically to earn the attention of networked audiences on large social media platforms. Via four case studies of Twitter discourse in the U.S., each of which presents a new method for understanding influencers' behaviors and organization, I seek to understand how to represent and visualize influencer communities; how to analyze the strategies that influencers pursue within these communities; and how to understand the effects these influencer communities have on political discourse in the United States. I conclude this dissertation with a collection of reflections, most of which concern how characteristics of influencer communities, such as their embeddedness in power relations and their relationship to the "mainstream," structure which strategies influencers within them take.

Acknowledgements

This work has been funded by University of Washington's The Center for an Informed Public, The John S and James L Knight Foundation (G-2019-58788), Craig Newmark Philanthropies, the William and Flora Hewlett Foundation (2019-9221), the University of Washington Population Health Initiative Pilot Grants, the Election Trust Initiative, the National Science Foundation (grants 2027792, 2120496, 1749815), and the Omidyar Network. For travel to present work contributing to my dissertation, I have been funded by the University of Washington's Center for an Informed Public and their Human Centered Design & Engineering program.

Otherwise, so much to acknowledge!

Thank you Mom and Dad for all your love and support and especially KIOK. Thank you Waurs, Emily, and Mike for being, in each a different way, my role models, and from age four onward giving me so many reasons to get excited around the holidays. Thank you Patty, David, David, Leigh, Chris, and Larry for so much of your support up and down the first three decades. Thanks to the boys for letting me ask them, like an old skeleton, what social media they paid attention to. Also for the wrestling matches, and for letting me relive my golden era of video games.

Thank you Emma and Kate for your mentorship, fearlessness, and humor. Thank you Jevin, Mako, Adrienne, and Joan for helping me one way or another finish this whole thing up.

Thank you to Jayashree, who taught me that most attempts to measure things are actually bullshit, and that achieving something good is as much a practical as statistical art. Thank you the rest of QTIM — Ken, James, Yi-Fen, Elizabeth, Kathi, Jay, — for mentoring me one way or another into becoming an academic. Thank you Yi-Fen for all the lunches when I first started in our lab. I often think about your stories about exploring the desert in between night shifts at the particle collider as a model for how I should do my own PhD.

Thank you to Izzi. As an academic peer, you're literally incomparable, and every conversation feels like I'm barely keeping up with you. But that doesn't matter at all; what's important is that

you've been my best friend since I've started this PhD. Thanks for Ibb and Obb, for every conversation, but especially for talking Shit, for always getting amazing gifts, for pizza and nuggets (thank Jackson for that too). I wouldn't remotely be the same person if I hadn't met you. If I win the mentorship award, your check is in the mail.

Thank you to Sarah, my closest collaborator on a nightmare project, and my first true friend at UW. Thanks for letting me mentor you for all of five minutes before you surpassed me in every respect. See you at the Base.

Thank you to Mallory for being as nightmarishly obsessed with this work as I am, and for trading so many articles and conversations during your short visit. Congrats on your own diss, I wanna be you when I grow up, etc. etc.

Thank you to Paul, Alex, and Lia. It goes without saying that you should be at the top of everyone's acknowledgement list. It was a joy working with you, and as someone who did your job before I came to the PhD, better you than me for all our sakes.

Thank you to Julia, Sawyer, and Michael. I am so in your debt for all the help you've given me at the CIP, like *everyone* in the CIP is. Julia and Sawyer, sorry I never learned how to book rooms.

Thank you to Danaë, Ethan, and Timnit for routing me, through kind responses to cold emails, towards UW and HCDE — I wouldn't have started a PhD without you, let alone enjoyed it. Thank you to Dana for talking me through so much of the PhD, and teaching me Python when I first came to Boston looking for a job. Congrats on being the next one through the gate.

Thank you to Melinda, whose refusal to help me with a SQL query on the first week of the PhD instilled in me a valuable ethic of self-reliance. You were my first mentor in the PhD program, and I've learned so much from you every time we've talked. I'm pretty sure 50% of the CIP's graduates will be working for you someday. Thank you to Tom, Himanshu, and especially Ahmer for your mentorship throughout the PhD. I had Tom and Ahmer's dissertations on tabs right next to mine the whole time I wrote this.

Thank you to my friends at the CIP and iSchool. Thank you to Zarine for having only correct opinions, especially at happy hour; thank you to Leo for having only correct opinions, especially on Prosser; and thank you both for Mario Kart and Smash and all your love and support. Thank you to Morgan for having one million ideas that I never helped on even once. Thank you to Ian for encouraging my ranting during the first EIP, and becoming a true friend afterwards. Thank you to Rachel, for being the CIP's greatest research and life role model, and always being willing to help anyone with a moment's notice. As I told Taylor, it's Rachel's world, and we're just living in it. Thank you to Maddy, Koko, and other Anna for your mentorship — I learned so much every time we talked. Thank you to Emily for being the potluck's biggest fan, which meant a lot to me. Thank you to Kayla and Sydney for letting me bug you whenever I wanted in Allen. Thanks to the rest of the CIP for being such a fun, weird place to work.

Thank you to the department. Thanks to Joice and McKane for such lovely research conversations, such lovely happy hours, for stealing my birthday party ideas, and for inspiring the purchase of the Nugget, which perhaps materially helped me finish this dissertation more than anyone else. Thanks to Caitie for being so much fun at all my potlucks and being a lovely friend. Thanks to Emma for being my get-out-of-here buddy, and being someone I felt like I could turn to for advice during the job search. Thanks to Tricia for the pizza recommendations, the dog Instagram story photos, and the opportunity for me to buy a mocha cookie crumble. Thanks to Julie for loving the east coast as much as I do, and for giving me the opportunity to buy a mocha cookie crumble. Thanks to Keri for helping me repair my laptop and making me feel like queen of the nerds for doing it. Thanks to Negin for always bringing something to the potluck, and particularly for bringing that lovely dessert, I want the recipe!! Thank you to Akeiylah for always being someone I can stop by in Sieg, and for having a karaoke birthday party to remember. Thanks John for getting out the door first, and sharing your wisdom along the way. Thank you to Beth for lending me the pod, which was certainly a learning experience. Thanks to Neilly, Jay, Lotus, Becca, Gabrielle, Andrea, and Meena for being such an amazing cohort for five years. Thanks to Taylor, Mariam,

Maya, Monica, René, Megan, Sehrish, Gauri, Swati, Trushaa, Alvina, Hel, Jules, Mads, and Annie for being awesome team members in Directed Research Groups.

Thank you to Bathhaus. Thank you Belle for feeling like a slice of the east coast from Idaho, and I think maybe just being my coolest friend, straight-up. Thank you Lydia for your ease and lightness in the house, and showing that you can be a literal maniac while still becoming a Doctor. Thank you Shard for changing my life completely, for being my best and deepest friend, and for making it in the knick of time for shot day. Thank you Kayla for being my first friend in Seattle, for having such faith in me and Bathhaus, and for stealing my trauma book. Thank you to Pixel for proving that some things are worth getting your arms scratched up for. Thank you to Gonzo, for showing that you get back what you put in.

Thank you to the Burke House. Thank you to Gene for being the Great Gatsby of Wallingford, and for bringing so many amazing people into my community through sheer will and joy. Thank you Sam for doing the same, through sheer will and gruffness. Thank you to Eric, Asela, and Laura for the long nights playing Blokus, and thank you all for frequently losing. Thank you to the Burke House Community Association, A.K.A. Ari and Cherrie, for being lights of community-building in a cloudy city. Thank you to Ben, who was born in the Burke House, and who has been my best friend, a bon vivant without compare, and someone I can always rely on.

Thank you to Vigilante Pet Rescue and their groupies. To Deianna, I have no idea why you trusted any of us at Burke House to move in, but I'm glad you did. Super proud of you and your cool new job all the time and you're my favorite friend to throw into any mix. To Jess, please pardon me when you're elected president of the United States. To Alex, please hook me up when you're elected CEO of TaskRabbit.

Thank you to the Zoup House. You made Boston seem like one of the best cities on earth when I stayed with you over the summer, and you all know how impressive that is. Thank you Nick for being such an amazing person to talk research with, and an amazing person to learn from, and for

letting TK hang out with me. Thank you Naomi for being so warm and inviting from day one of my visit — truly, the sort of person that starts a co-op. I learned so much from our conversations, and your work is so cool that it makes me re-evaluate my entire life.

Thank you to Club Lola. That you Jackson for having a better class analysis than pretty much every UW PhD, for your generosity day-in and day-out, for putting Seattle nightlife on your back, and for putting up with my shit in general. Thank you Lidia for being the definition of an icon, while also being the definition of kindness, and sharing my love of picking people up from the airport. Lola, thank you for teaching me how much time it takes to do one thing well, and for reaffirming that the meaning of life is to hang out with your friends.

Thanks to the Sapphics — Amanda, Lauren, Victoria, Katherine, and Rachel — for chaos, friendship, and one of the best secret Santas I've ever been a part of.

Thank you to Emily and Leigh for being my best friends, for liking Fireball and birds respectively, and for being my most constant supports over the last 10 years. Thank you to Lev, Caleb, and Jonny for being my best friends, for enduring what Caleb recently referred to as my "podcast," and for spending so much time talking about those three red lights. Thank you Helen for being my best friend, and the smartest and most generous person I know, and for birthing Fig into this world.

Thank you to Andrew for being my best friend and providing me with the worst sublet of my life. Thank you David for being my best friend and being such a joy to visit in literally any context. Thank you Nick for being my best friend, and being one of the first people to push and challenge me. Thank you Will for being my best friend, for having my exact sense of humor, for always supporting me, for being a model for how to pursue your dream, and for Burritos Pizzeria. Thanks Snake for nothing, and for being my best friend, and for always reaching out — I really appreciate it.

Thank you to Rachel, my best friend. You were so important to me these last five years, and the five years before that, and the next fifty years after this.

Thank you to Robin and Dessie for tunes and broths, and for starting my move to Seattle off with the correct road trip. Thank you to Hailey for not even pretending a dissertation was something to care about, and for being my best friend. Thanks Nick for bringing the third skull, and for a really interesting conversation right before defending this dissertation. Thank you to Margot for crashing through life with me at the right time, and continuing to be an amazing friend a state or so away. Thanks to Mara and Trevor for basically always pushing me be less dumb, and for tolerating the burrata. Thank you to Quiet Michael for always taking me completely seriously, and for all our conversations.

Thank you to Meagan for crashing into my life so late into the PhD. You're such a source of great advice, big joy, and dancing. Looking forward to seeing your octopus costume. Congrats on getting me to stay in Seattle for a little bit longer. Thank you to Mariya and Jeffrey for being my first friends in Seattle — it's a gift to welcome someone into their new home.

Thank you to Elowyn and the rest of the Lindy crew, for showing that there's always another loving community waiting right around the corner, even in year five. To Elowyn specifically, you have no idea what a boon you were for co-working productivity.

Thanks to S, who introduced me to science and technology studies way back when, and was one of the only people to take time off to join me on my ridiculous cross-country roadtrips. Thanks for the visit while I was writing my dissertation, do stop by again.

And to Madeline, for the tiny beverages and well-plated lunches, for the co-working sessions that were and weren't, for listening to my rants and thanking me for sharing, for teaching me so many things while still treating our lives like one big game. You were my biggest help in finishing this dissertation, and my greatest reason to.

Contents

List of Figures

7.1 **A visualization of the cited literature of scientists and science communicators with respect to mask-wearing and COVID-19**. Papers that tend to be cited together by A) scientists and B) science communicators discussing masks and COVID-19. Scientists tend to cite groups of papers by topic or field. In A), the green cluster represents papers concerned with the physical properties of mask materials, the orange cluster represents epidemiological papers on masks effectiveness against COVID-19, and the purple cluster represents older public health papers on the effectiveness of masks in hospitals and households for COVID-19 and other respiratory illness. In B), meanwhile, science communicators cite largely the same papers as scientists, but segregate their citations by stance: supporting the use of masks in blue, and opposing them in red. Three papers (nodes) are labeled in both clusters: 1) (Leung et al., 2020), the most shared paper in both datasets, 2) (Rengasamy, Eimer, and Shaffer, 2010), a fluid physics paper from 2010 shared in the anti-mask perceived consensus, and (MacIntyre et al., 2015), a 2015 paper with a negative finding on masks whose lead author has since publicly supported mask-wearing for COVID-19. Larger nodes in the scientist network (A) are cited by more highly-cited papers on masks, while larger nodes in the science communication network (B) are cited by more frequently-shared Twitter URLs in a dataset of people arguing about masks. 181

List of Tables

Dedicated to Mom. KIOK!

Chapter 1

Introduction

The United States (U.S.) is in the midst of a paradigm shift in who creates news. It is widely known that internet has in the past few decades displaced television and print newspapers as the primary medium where news is consumed (Shearer, 2021). However, the past two decades have also seen a shift in *who* is communicating the news. People consuming digital news, and especially young people, are now less likely to get their news from journalists or television broadcasters, and more likely to receive it from a nebulous figure commonly referred to as the influencer (Robb, 2020; Newman et al., 2023). People are not only paying more attention to influencers, but there has seemingly been a remarkable growth in the type of and number of influencers in the last decade, enabled both by the maturation of new, massively-networked social media platforms and the professionalization of influencing as an industry unto itself (Duffy, 2017; Abidin, 2018; Hund, 2023).

The rise of the influencer in the creation of news, and U.S. political discourse more broadly, has come with its fair share of public anxiety. Particular attention has been paid to how influencer practices have, in the last decade, become well-integrated into the distribution of U.S. far-right and conspiracist content online (Lewis, 2020; Ma, 2021; Moran, Swan, and Agajanian, 2024). Some have worried about the corresponding spread of misinformation and radicalizing content from these "superspreaders," and its potential to mutate into offline political action (Kennedy et al.,

2022; Karell et al., 2023). However, the focus on the worst excesses of influencers simply reflects a more general reality, which is that the logic of influencing that had once been restricted to "commercial" discourse like fashion and gaming has now broadened into self-conscious political discourse writ large. Political communication now happens, for good and for ill, via the mediation of influencers and the forces that govern them.

In this dissertation, I use mixed-methods to analyze a series of case studies of U.S. political discourse networks to better understand the structure, tactics, and effects of our new era of influencer media. After conducting a literature review of previous work on influencers in political communication, media studies, and human-computer interaction, I propose flexible ways to represent both influencers and what I call influencer communities in social media networks. Influencer communities are large assemblages of influencers working with and against each other to earn the attention of networked audiences on large social media platforms. Via four case studies of Twitter discourse in the U.S., each of which presents a new method for understanding influencers' behaviors and organization, I seek to understand how to represent, visualize, and analyze influencer communities; how to understand the strategies that influencers pursue within these communities, and how to understand the effects these influencer communities have on political discourse in the United States. I conclude this dissertation with a collection of reflections, most of which concern how power relations and proximity to the "mainstream" structure the behavior of influencers and influencer communities.

1.1 Motivations for this Dissertation

As a preamble, I describe three motivations that led me to structure this dissertation as a mixed-methods study around U.S. influencer communities.

First, I am particularly motivated to study the influencer community, as distinct from individual

influencers or influencers as a static type. Ethnographic researchers in cultural and media studies have produced excellent studies of the behaviors and motivations of individual influencers at the micro-scale, while experimentalist researchers in political communication have similarly analyzed the general behavior of influencers as a type in the macro-scale. My goal here is to study influencers at a meso-scale, where aggregates of influencers come together in stable yet shifting communities defined by shared audiences and discourse contexts. I do not specifically study individual interactions between influencers, or between influencers and their fans, but rather the nature and effects of aggregate interactions between influencers and audiences. I take it as given that different communities have different norms of behaviors and stances towards other communities, forming a much larger network of influencer communities with varied interaction patterns structured by norms, values, and power. My aim, then is to forefront this way of analyzing online influencer dynamics and to demonstrate its value in generating empirical insight about particular media ecosystems.

Second, I am particularly concerned with the task of measurement *and* representation of influencers and their communities, and I view these two aims as fundamentally inextricable from each other. The definition of an influencer, like all social roles, is constantly being renegotiated — and particularly so in this last decade of their irrepressible growth. My proposal to measure the activities of influencers in political media necessarily entails defining representations of these influencers and their aggregates. My goals, accordingly, are not to create measurement tools which accurately capture a phenomena compared to an already-known ground-truth, but rather to propose exploratory tools that aid in the process of renegotiating the influencer concept itself. The evaluation of these methods comes not from fidelity, but from the extent to which they clarify conceptual boundaries, reveal useful empirical findings to practitioners, and generate further productive research questions about the nature of influencers in our social world. My aim thus is to provide new ways of structuring social data which can provide a productive foundation for future studies of influencers.

Finally, I stress that this is a case study of *United States* political media. I do not seek to ask research questions whose answers will generalize to all cultural contexts of the internet, because I take it as given that such generalizations are impossible. Social processes cannot be extricated from the histories and cultural contexts where they occur, and so it goes for the study of influencers in the United States. Each of these chapters uses the case study format, where measurement tools are proposed, applied to large datasets focused on the U.S. digital context, and then analyzed for what insights they can (and cannot) reveal about these contexts. You will read in this dissertation extensive results related to how the dynamics of influencer communities manifest themselves in the particular instance of U.S. political culture, and the presumption that these manifestations are, in and of themselves, valuable empirical findings. My motivation here, then, is to participate in the ongoing reclamation of the mixed-methods, historical case study as a productive way to study online phenomena.

1.2 Research Questions

This dissertation seeks to answer three research questions with respect to influencer communities, which correspond to three general themes: structure, strategy, and effects.

RQ1 (*Structure*): What methods can we use to characterize the members, boundaries, and interrelations of different influencer communities? (Chapters 4, 5, 6, 7)

RQ2 (*Strategy*): How can we characterize the strategies that influencers use in the aggregate to attract sustained attention, with respect to their own communities and their peer communities? (Chapters 4, 6, 7)

RQ3 (*Effects*): How can we measure what influencer communities do, and particularly how their activities structure online political discourse? (Chapters 6, 7)

As indicated, these research questions are distributed across four studies, with different aspects of studies reflecting on individual research questions. Other aspects of these questions are addressed

in the Background, Methodology, and Discussion chapters, which take broader views than individual studies can provide.

1.3 Structure of this Dissertation

The structure of this dissertation follows below. The remainder of this chapter outside of this section will discuss the contributions of prior work to this dissertation.

Chapter 2 (Background). This chapter conducts a relatively brief literature review on the influencer concept from the perspective of political communication and media studies. This background section contextualizes some of the choices made in the representation of influencers in subsequent chapters.

Chapter 3 (Methodology). This chapter provides a brief literature review of the methods used in this study, namely network science and the mixed-methods case study. It also describes the nature of the social media datasets analyzed in this dissertation, and how these methods may generalize beyond these datasets. I then provide working representations for the concepts of "influencer" and "influencer community" for use in subsequent chapters, derived from the theoretical literature summarized in Chapter 2. I conclude with a brief discussion of my positionality with respect to this work, and how it may affect my findings.

Chapter 4 (Study 1: Coengagement Networks for the 2020 U.S. Election). This chapter introduces the fundamental method for the first three chapters, and which is closely related to the method in the final chapter. It provides a basic tool for representing influencer communities (RQ1), and provides empirical findings in the context of the 2020 U.S. presidential elections using network metrics. It discusses the use of *triangulation* in the visualization of online communities, and contrasts the coengagement graph with ways of visualizing networks.

Chapter 5 (Study 2: On Measuring Change in Networked Publics). This chapter uses the framework of the coengagement graph to analyze an influencer community longitudinally, in this case

the same community identified in Study 1 projected forward into the 2022 presidential election. It proposes basic methods for measuring how influencer communities change (RQ1), and along which axes, and reflects on how change in political influencer communities is a topic of fascination for larger publics.

Chapter 6 (Study 3: Negativity in U.S. Influencer Culture). This chapter uses the coengagement graph method combined with natural language processing algorithms to look at structures of negative interactions between political influencer communities (RQ1). It particularly inspects why influencers use negativity, the effectiveness of negativity as a method for gaining followers, and the targets or negativity in these influencer communities (RQ2). This uses a longitudinal dataset of influencers tracked from 2021-2023.

Chapter 7 (Study 4: Dueling Consensuses). This last chapter concerns itself with a new dataset of largely health and science influencers, some of whom are mainstream (e.g. science journalists) and some whom are conspiracy theorists. Moving away from social media platforms, we show how the websites they create get motivated as "evidence" which contributes to the appearance of separate, dueling scientific consensuses (RQ1). We show how these influencers use negativity to designate anti-consensuses among their audiences, and draw boundaries between science and not-science (RQ2). We discuss the implication of consensuses being the outputs of influencer communities: portable informational products that can be put to many uses by their audiences (RQ3).

Chapter 8 (Discussion). This chapter summarizes contributions made in the four studies of this dissertation, and reflects on future opportunities for research on influencer communities.

1.4 Integration of prior work

Studies 1 and 4 in this dissertation have been previously published in a peer-reviewed conference and journal respectively (Beers et al., 2023a; Beers et al., 2023b). Study 2 is, as of writing in May

2024, under revision at a peer-reviewed journal, and its text in this dissertation has only been slightly modified from its original submission. Study 3 is being prepared for submission at a peer-reviewed journal. An earlier version of Study 3 was presented at a lightly-reviewed conference, and two conference submissions incorporating aspects different of that study's work have been submitted as of writing (Beers, 2023). Some text from these other submissions are incorporated into Study 3.

For Studies 1 and 4, I have not intentionally modified the text or figures from their original form, and in each case their peer-reviewed formats should be considered their "final" forms. I include a short preamble before each study to connect their contributions to the larger structure of this dissertation. Cited work is combined into a shared bibliography at the end of this dissertation.

Studies 1, 3, and 4 were created with collaborators, and within the study chapters, while Study 2 is solo-authored. In each case, I am primarily responsible for, at least, the conceptualization and writing of these papers. Accordingly, in the Methodology chapter, I use the first-person "I," although the selection of many methods benefited from discussions with my co-authors and colleagues. In the published works, I am first author, and am assumed to be first author in the latter two studies in preparation for publication.

Some materials in Chapter 2, Background, were adapted from previous unpublished work submitted for my candidacy exam and thesis proposal.

Chapter 2

Background

The term "influencer" is of relatively recent coinage, derived more from popular discourse than academic research. In response to the rising popularity of this term, different disciplines have had to map different existing concepts into the new category of influencer. I review here the simultaneous evolution of the term "influencer" from two prior concepts: opinion leaders in political communication, and celebrities in media studies.

We also see the study of influencers emerge in some other fields that I do not review here. The largest absence is the field of marketing, which primarily focuses on the effectiveness of influencers for generating revenue for advertising companies, and the characteristics of influencers which drive audience engagement (Vrontis et al., 2021). Given the relatively separate development and goals of marketing research from the three studied in this chapter, I leave the full integration of this research lineage to future work. Studies of influencers and social influence also increasingly penetrate the fields of sociology (e.g. Rohlinger et al., 2023; Karell et al., 2023) and social psychology, inasmuch as it is increasingly difficult to study media dynamics without encountering influencers. However, the topic is currently less well-theorized in both fields, and largely borrows terms from the two profiled here.

2.1 Political Communication

I begin with the influencer concept in the field of political communication. Political communication considers communication by elected politicians, addressed to politicians, or about them and their activities generally (McNair, 2017). It can, at times, be difficult to draw boundaries between political communication and other fields, due to the significant extent of shared terminology and shared concerns it has with fields like cultural and media studies. I use this background to (1) explain certain normative motivations for this dissertation, namely the potential for influencers to sow so-called misinformation and polarization, and (2) provide productive contrasts between highly developed communication theories and more modern ethnographic research centering on "influencers."

2.1.1 From Opinion Leaders to Influencers

Political communication contains one of the earliest theoretical lineages for what would become the influencer: the opinion leader. Research on influencers in other fields typically focus on their practices: their strategies for growing and maintaining audiences, and for translating attention into capital. By contrast, the definition of opinion leaders was more ontological: they were, simply, people who had by one method or another achieved influence upon their peers' opinions. Taking the existence of opinion leaders as a given, scholars of political communication sought to explain the effects opinion leaders had on the interpretation of political media and political discourse more generally.

Opinion leaders as such were first defined and elaborated by sociologists Katz and Lazarsfeld (1955). They were hypothetical people who served as a bridge between media outlets, which they presumably consumed heavily, and the broader, disinterested public. This so-called two-step flow of information (from media to opinion leaders, and from opinion leaders to the public) challenged the dominant one-step flow of information, where mass media simply provided beliefs to a pliable and receptive public. Both models, however, had worrisome implications for the

long-term survival of democracy, which at the time of was somewhat in doubt due to the rise of fascism during World War II (Gitlin, 1978; Pickard, 2021). On the one hand, a one-step flow which privileged mass media suggested that if mass media were to be co-opted by, say, an authoritarian state, it could be an incredibly effective tool of social control. On the other hand, a filtering step of "ordinary" people interpreting mass media for their audiences suggested that existing social networks of friendship and affiliation were resilient to interference from mass media, which could in theory contain "objective" information and a shared basis for reality.

The next few decades of opinion leader research mostly consisted of a fierce debate attempting to confirm or deny opinion leaders' existence in the mass media environment, and quantify the extent of their influence on their audiences, if any. Weimann (1991) synthesized many of the debates over the potential existence of opinion leaders, and suggested a new term: "influentials." Weimann's term presaged a cultural shift that would occur with the advent of the internet. Online networked communities, political or otherwise, started to create people popularly perceived as influential, whether political communication scholars believed they were or not. By the end of the 2000s, news outlets such as the New York Times began to legitimize a word that was already circulating to name these people: influencers (Stone, 2009).

Some researchers, such as Bakshy et al. (2011), adapted the existing opinion leadership debate to this new terminology, testing for the existence and relative effects of "influencers." However, other researchers took advantage of the unique affordances of the internet to propose new mechanisms supporting or opposing the two-step-flow theory. Some researchers had renewed enthusiasm for the one-step flow specifically or more generally the dominance of industrial media, arguing that personalized algorithms were a new, more powerful broadcast media that injected content directly into internet users without the mediation of opinion leaders (Bennett and Manheim, 2006; Weigel, 2022). Other researchers endorsed the two-step flow, arguing that the unique affordances of the internet, combined with the emergence of the online influencer, still allowed for significant interpersonal mediation (Bennett and Iyengar, 2008; Choi, 2015; Hilbert et al., 2017). While old

debates persisted, new data sources and new assumptions provocatively changed their terms.

Modern panics over misinformation, explored more thoroughly in subsequent sections, borrowed anxieties from the one-step-two-step debate, also called the media effects debate. If institutional sources of truth — regarding, say, the effectiveness of COVID-19 vaccines — were unable to relay their message to the public via mass media without corruption by intermediate partisan influencers, then misinformation would surely result (Bennett and Iyengar, 2008). The impetus to study influencers then partly reflects an anxiety about their responsibility to communicate a shared reality which was supposedly broadcasted by traditional mass media sources — a responsibility that surely some are shirking for self-interested motivations.

2.1.2 Publics and their Influencers

Scholars of political communication also needed a term to describe the aggregate activities and effects of opinion leaders (if they existed) and mass media sources. The theory of publics, sometimes collectively called the public sphere, proved useful for situating opinion leaders and later influencers. Most definitions of publics have their root in the Habermasian public sphere, generally defined as a mediated space where a nation's private citizens come together to deliberate upon "public" issues in a structured and rational manner, with the goal of informing governing priorities (Habermas, 1991). The public sphere was originally conceived as an ideal that was, at the time of writing, already being corrupted. This hoped-for space for private individuals to deliberate over ideas in groups was rarely actually achieved, Habermas worried, due to the intercession of market and state influence as mediated by communication technologies.

Habermas' description of the original, ideal public sphere was originally limited the activities of a small, elite group of men who had access to newspaper editorials and discussion salons. Perhaps unintentionally, this description was prescient: the most popular contemporary political influencers in our studies are also a small, disproportionately powerful group of mostly white

men[1]. But more broadly, this description failed to account for the fact that group deliberation of political issues happens in many ways and in many locations, and that different deliberating groups interact with each other in ways that are conditioned by power. Fraser (1990) first proposed the concept of subaltern counterpublics, which were alternate spheres of discussion where marginalized groups could discuss issues and eventually raise them to a distinct, dominant public sphere. Squires (2002) elaborated on counterpublics in the context of U.S. online Black communities, proposing new types of publics (*enclave*, *satellite*) which were not necessarily engaged in direct debate with the main public. Recent scholarship has shown have subaltern counterpublics not only negotiate with the dominant sphere, but also with each other, in what Gutiérrez (2022) has called competing ethnoracialized counterpublics or Billard (2023, ch. 6) the "politics of flows." While most scholarship on alternative publics focused on marginalized communities, some have implied that the same frameworks extend to far-right and authoritarian communities (Downey and Fenton, 2003, p. 197), recently dubbed by Jackson and Kreiss (2023) as "defensive publics" for their defense of rather than challenge to status quo inequalities.

Scholars of publics now generally take it for granted that group deliberation occurs not in one single public sphere, nor even in a hub-and-spoke system with a dominant sphere at its center, but as an interdependent network of publics. This emergent concept of the network of publics is what motivates this dissertation's study of influencer communities at the meso-scale. Many of the nodes in this network of publics consist of communities of political (or politicized) influencers, who do the work of negotiating community priorities and values. These influencers do not simply advocate against an intransigent dominant public, but engage with each other from shared (but

[1] In Study 1, Case 1, the top 10 most central pro-Trump influencers are white men, while 8/10 pro-Biden influencers are white men or organizations led by them. In Case 2, 9/10 pro-Trump influencers are; 7/10 pro-Biden are; 6/10 pro-socialist are. In Case 3, 8/10 pro-Trump influencers are, 7/10 pro-Biden are, 7/10 pro-socialist are, with followback accounts' identities indeterminate. In Study 2, the 10 most-retweeted right-leaning users are white men or organizations led by white men in 2020, and again 8/10 are white men in 2022. For left-leaning users, 6/10 of the most-retweeted users are in 2020, and 7/10 in 2022. In Study 3, 8/10 of the most central pro-Trump influencers are, 7/10 of the most central progressive mainstream influencers, and 6/10 radical socialist influencers. Study 4 is more difficult to quantify, with many of the most-shared URLs being authored anonymously and by large organizations, but generally shows a more even gender distribution than the election-focused case.

unique) positions of marginality or defensiveness. The concept of opinion leaders within publics is a productive point of comparison for the concept of influencer communities, which I develop further in the Chapter 3 (Methodology).

It should also be noted that Habermas was initially concerned with the effects of mass media technologies and industries, such as those supporting broadcast television, on the democratic potential of the public sphere. Internet scholars have taken up this call to study the transformative effects of media technologies on publics, theorizing a range of new publics including the blogosphere (Farrell and Drezner, 2008), the Twittersphere (Bruns et al., 2017), Twitter publics (McKelvey, DiGrazia, and Rojas, 2014), and networked publics (Boyd, 2007; Varnelis, 2008). boyd's networked public, which was defined as both the internet-mediated space upon which publics form and the resulting publics themselves, became a leading term, and in general "publics" became a scholarly shorthand for referencing hierarchical social media communities. Scholarship on networked publics has often focused on counterpublics, in some cases arguing that social media design affordances provides new opportunities for counterpublic engagement with dominant publics (Jackson and Foucault Welles, 2015), and in some cases arguing that internet surveillance reduces the effectiveness of certain counterpublics (Salehi et al., 2023).

The study of publics in online communities has shown that technology, and the industrial and cultural phenomena that accompany them, structure the normative effectiveness of publics. It is no doubt expected then that the rise of influencers, who derive their unprecedented popularity from the unique affordances of social media and internet marketing industries (Hund, 2023, ch. 3), would also change the effectiveness of the networked public sphere. This dissertation thus develops tools which can characterize new influencer-dominated publics, with hope of evaluating influencers' effects on the normative goals of publics for rational policy deliberation.

2.1.3 Influencing Audiences

The normative concept of the public has consistently been contrasted with the primarily descriptive concept of the *audience* (Livingstone, 2005; Butsch, 2011). In contrast to the public, whose media and members were supposed to meet certain standards in order to support a functioning democracy, audiences were people who merely received media messages on any topic. Members of the audience were sometimes critiqued as passive, atomized individuals, absorbing messages from a consumerist media which primarily entertains and enriches, rather than informs and democratizes (Morley, 1993; Livingstone, 2005). However, many scholars, particularly those more oriented in media and cultural studies, have advocated for the agency and power of audiences, describing the deeply social and interpretive work that audiences perform on media (Hall, 1980). The internet era, with its increased emphasis in audience ("user") interaction in media, particularly stressed the active role that audiences could take, and indeed questioned the stability of media/audience distinctions entirely (Bruns, 2007; Green and Jenkins, 2011). Understanding the nature and capacities of audiences is critical to understanding what drives influencers, as influencers are often taken to be simultaneously manipulating contemporary audiences, and in thrall to them.

Audiences are groups of people viewing media, shaped both by available media technologies and prevailing norms about how media should be experienced. Butsch (2011) describes how common perceptions of U.S. audiences evolved from the raucous, distracted scenes of socialization of 1800s theatre halls, to the family-based scenes of alternately devoted and intermittent attention towards radio and television, to the finally isolated-yet-networked spaces of users of the internet (see also Butsch, 2000). Per Butsch, changes in the physical presence of and interaction between audience members led to idiosyncratic changes in how they were normatively perceived. Communal, physically-situated audiences could be productive in, say, a policy debate, and harmful in a distracted mob. Mass, global/national audiences could inspire hope for the prospects of widely-informed, democratic public in one formation, and fear that their mass-ness would lead

to counter-productive passivity and media domination in another. Butsch and Livingstone (2014, p. 9) argue that whether audiences are judged as more or less like publics has as much to do with the marginalized (or dominant) status of their members as it has to do with any other empirical qualities of these audiences.

Paralleling the one-or-two-step communication flow debate, there has been a debate over the extent to which audiences are "passive" (in the thrall of media messages) or active (re-interpreting, re-creating, or encoding media messages). Hall (1980) proposed an alternative relationship between media and audiences, where media send messages whose effects are contingent, rather than controlled. Media encodes messages in commonly understood signs and symbols, while audiences decode messages according to their experience and stances toward that media. While it is important to note that media still retain some agency to structure messages' effects (Morley, 1993), subsequent scholars have used similar frameworks to Hall to argue that audiences are resilient and active generators of meaning from received media messages. hooks (1992, p. 116) suggests that audience spectating can be a powerful form of resistance in and of itself, praising the power of the *oppositional gaze* in Black female spectators:

> Spaces of agency exist for black people, wherein we can both interrogate the gaze of the Other but also look back, and at one another, naming what we see. The "gaze" has been and is a site of resistance for colonized black people globally. Subordinates in relations of power learn experientially that there is a critical gaze, one that "looks" to document, one that is oppositional. In resistance struggle, the power of the dominated to assert agency by claiming and cultivating "awareness" politicizes "looking" relations — one learns to look a certain way in order to resist.

Internet technology has only increased opportunities for audience agency and resistance, and increased the potential for this agency to occur in large, networked settings. Baym (1999) documents how Usenet enabled the formation of bustling fan communities for soap opera viewers, where the focus was as much on interpersonal connection between fans as it was on deconstructing media

messages.[2] Updating hooks (1992), internet scholars have particularly analyzed how marginalized communities have continued to challenge and build from hegemonic narratives, albeit within the structuring power of mass media messages and platform technologies. Steele (2021) describes certain online spaces for Black social media users as a "virtual beauty shop" where Black feminist discourse is constructed often through the consumption of media, while Bailey (2021) describes the "digital alchemy" of transforming media *misogynoir* into content which supports Black feminist community- and identity-building. Oppositional audience communities can also be formed for less obvious prosocial and counter-hegemonic purpose, as literature on antifandom and "hate-sharing" makes clear (Click, 2019; Gray, 2021; Weigel, 2023; Weigel and Gitomer, 2024). As we have moved into the influencer era, studies of antifandom have particularly focused on hate targetted towards particular influencers, as I study in Study 3 (Chapter 6, Chin, 2019; McRae, 2020; Duffy, Miltner, and Wahlstedt, 2022).

Media generally and influencers specifically have always been guided by the interests of audiences, or at least their perception of them. Marketers have been developing increasingly sophisticated and diverse methods for nearly a century to measure the activities of audiences, to measure the effects of both advertisements and more traditional media (Ang, 1991; Napoli, 2010; Balnaves, O'Regan, and Goldsmith, 2011; Nelson and Webster, 2016). Contemporary scholars of social media have noted how the daily practice of social media, for users and media-makers alike, requires certain processes of "imagining" an audience, whether directly through metrics or other qualitative processes (Litt, 2012; Marwick and Boyd, 2011). As I document more in the following section, contemporary influencers use a variety of marketing tools combined with their own intuitions to track the engagement of their imagined audiences (Hund, 2023). Webster (2011) cites Giddens (1984) to argue that the increasingly metricized and predictive online environment leads to structuration between media systems and audiences: the mutual, iterative co-construction of structure

[2] Baym's work itself drew on a then-emerging, and since well-developed field of fandom studies. Jenkins (1992) perform one of the canonical studies of fandom within the field of cultural and media studies, with Duffett (2013) providing an updated, general survey.

through the intermediary of likes, share, comments, and other metrics. The history of media interpretation of the audience reveals that media constructs itself through its vision of its audiences, and thus that an understanding of media entails a understanding of that vision.

This brief review of audience theory informs the primary methods of this dissertation, which attempt to represent imagined influencer communities through the intermediary of a shared audience. Via the methods I propose, influencer communities in some sense *are* their audiences, whose attention and/or engagement is represented by various direct and derived social metrics from Twitter. Audience theory suggests that it would be a mistake to confuse these simple metrics with audiences themselves, taking audience members as simple voters or thermometers of public opinion. In the first place, audiences are always imagined by one party or another, making related concepts, such as public opinion or attention, always also imagined. But equally important is that regardless of how they have been imagined, audiences have always displayed a much more varied range of interpretive activities than simple endorsement or disapproval. Mere metrics cannot encapsulate the extent to which audiences are *active*, and creating and expressing meaning and feedback which exceed the constricted forms of likes and shares. Putative influencers do not control the effects of their messages, but rather craft and them submit them for uptake for a variety of audiences — even audiences they might not support, as I document in Study 3 (Chapter 6). As such, the methods proposed in this dissertation are attempts to develop stable proxies for a complex and irrepressible galaxy of audience forms and behaviors.

2.2 Media Studies

I continue with a brief theoretical history of the influencer from the perspective of media and cultural studies, and a selected historical review of influencer-like figures in U.S. political media. It is important to first note that a theoretical history of influencers from a U.S. perspective is not a history of influencer-like concepts globally. Abidin (2018, p. 2) notes, for example, how Japan has a

highly-elaborated celebrity culture that does not easily fit into pre-existing influencer concepts developed in U.S. contexts. The following review of influencers thus particularly serves the purpose of developing a case study of U.S. politics, but is incomplete as a total history of influencer-like concepts. The selected review of historical "proto-influencers" serves to contextualize that many of the media communities which are described in this dissertation, and highlight how they may or may not differ from pre-internet media manifestations.

2.2.1 Proto-Influencers in U.S. Media

Almost all writing about new technologies is a struggle to understand which features of a system are new and contingent on emerging technologies, and which are old, but perhaps defamiliarized through their mediation in new contexts. In light of this, I begin this section with a brief historical sampling of influencer-like figures in the U.S. Given the primacy of far-left and far-right publics to the case studies in this dissertation, I pay particular attention to media elites in these movements.

Benkler, Faris, and Roberts (2018, chap. 11) located the direct antecedents of particularly far-right social media influencers in radio, particularly after the repeal of the Fairness Doctrine in the late 1980s. Radio enabled well-resourced individuals to acquire for themselves the reach of traditional mass media formats, like newspapers or television. From the early 1900s to the present day, these radio hosts eschewed "objective" news reporting for partisan and yet also personal, funny, and endearing dialogue. While the Fairness Doctrine, which mandated equal coverage between opposing views on controversial political issues, hampered the activities of some far-right radio hosts, some survived and prospered even in the strongest government regulation regimes (Matzko, 2020; Brooke, 2022). While the contemporary form of right-wing internet influencer is most recognizable in radio, antecedents could also be found in Jim Crow newspaper columnists (McRae, 2018, ch. 4), and much of radio's partisanship and personality was also ported to broadcast television in the 1990s (Hemmer, 2022). Regardless of the era, far-right activists and entrepreneurs had always used media technologies to advocate for ideologies and styles of political communication which

were less tolerated in mainstream media institutions.

Far-right demagogues were not the only actors in the U.S. to use media technologies to circumvent mainstream institutional media. Progressive and radical groups have always created their own alternatives to mainstream media, and used whichever media technologies were available to them to broadcast their messages. These media efforts were sometimes much-needed correctives to mainstream reporting flaws, such as those provided by Black newspapers during the Civil Rights Era (Michaeli, 2016), those provided by online messaging boards during the early years of the HIV/AIDS pandemic (Cifor and McDonald, 2023), and those provided by progressive radio hosts on a range of topics on programs like Democracy Now! (Goodman, 2016). Popular histories of radical and progressive proto-influencers in the demagogic mode are more rare, but doubtless such figures exist in certain contexts, as they do on the internet today (Harris, Foxman, and Partin, 2023). And while left-leaning alternative mass media often hewed more closely to mainstream norms of objective and fact-based reporting than far-right media (Benkler, Faris, and Roberts, 2018), it is to this day not immune to spreading political misinformation (Beers, Wilson, and Starbird, 2022).

We can take from the history of influencer-like figures that media technologies provide tools for contesting dominant public spheres, both in frightening and encouraging ways. Technologies with lower barriers to entry allowed for competitors without institutional power, legitimacy, or capital to create their own messages that interpreted and contested the frames deployed by mainstream media outlets. As contemporary ethnographers document, internet media has drastically reduced the barriers to entry for creating alternative mass media, creating opportunities for both personal profiteering and platform-driven exploitation in the age of the influencer.

2.2.2 Microcelebrity on the Internet

Most of the historical work on proto-influencers in the previous section was driven by questions familiar to scholars of political communication. Many asked, for example: were these insurgent

media figures contributing to democratic deliberation, or eroding it? By contrast, much of the recent history ethnographic research in media studies has generally sidestepped the question of democratic deliberation directly, and instead seen influencers as laborers in a burgeoning and increasingly exploitative economic system of *attention*.

The study of influencers in media studies has its roots in the study of celebrity (Rojek, 2001; Marwick, 2010). While celebrities were studied from a variety of perspectives, Rojek (2001) describes celebrities as cultural commodities attached to real people; fictions created and maintained by large teams of culture workers (publicists, marketers, etc.), deployed through the mass media, and interacted with only parasocially, i.e. non-personally. He and many other celebrity scholars wrote at a time when the internet was both changing the nature of mass media, the nature and potential of parasocial interaction, *and* necessarily the nature of culture workers. Inevitably, the nature of celebrity would change with it.

As scholars began to update concepts of celebrity for the internet era, a new term emerged: microcelebrity (Senft, 2008). Doubtless, traditional celebrities like Hollywood celebrities had their practices changed by the advent of the internet and social media (Marwick and boyd, 2011). For example, the two-way nature of social media gave in theory if not often in practice a way for celebrities to break down and reconstruct the strictly parasocial relationship most of them had with their fans, potentially helping build and maintain attention. More significant, however, was social media's ability to lower the barriers for entry to the practice of exposing and commodifying one's image to audiences — to let *anyone* be a celebrity. To be clear, this does not mean that celebrity had been "democratized," as traditional celebrities with pre-existing capital still disproportionately reaped the benefits of social media's capacity for personalized sharing (Marwick and boyd, 2011). However, these lowered barriers opened the floodgates for people to *attempt* to be celebrities, and from this new audience of would-be celebrities — microcelebrities — came the kernel of a community of practice which would begin to shape the norms of contemporary influencer culture.

2.2.3 The Industrialization of Influence

It was not long after the advent of the micro-celebrity that scholars began to notice ominous economic trends among practitioners of micro-celebrities. Hund (2023, chap. 1) most directly describes how high unemployment caused by the 2008 U.S. recession pushed people, and particularly women, to try and recuperate income through the practice of microcelebrity. Senft, the pre-recession originator of the term, herself noted years later how a disaffection with reduced employment opportunities, combined with the massive expansion of personal media technologies, was increasing the appeal of micro-celebrity practices (Senft, 2013, "The Branded Self Online: The Paradox of Late Capitalism"). Duffy, in the aptly-named "(Not) Getting Paid to Do What You Love" and other work, described how despite a widespread perception of micro-celebrity practices as a labor of love, it was self-consciously seen by its practitioners as a way to supplement income streams during a time of economic precarity (Duffy, 2017; Duffy et al., 2021). She called their work *aspirational labor*, i.e. a precarious labor with high up-front time and effort costs that may never generate income.

Abidin's work in the mid-2010s on internet celebrity serves as a bridge from the prior concept of microcelebrity to the contemporary concept of the influencer. Abidin (2018, p. 14-15) theorized influencers as a subtype of a broader concept of internet celebrity that had several critical differences from micro-celebrity, including: (1) a potential for exceeding the reach of traditional or even national media (not "micro"), (2) having cross-platform instead of single-platform audiences, and (3) consequently having broad cultural and economic impacts befitting their increased status. As Hund (2023, ch. 3) and Abidin (2018, ch. 4) both describe, influencers co-evolved with marketing infrastructures designed to support and exploit them, creating an entire influencer industry centered around translating influencers' ability to garner attention into, for example, purchases of the commodities. This development constituted, per Hund (2023, p. 10), the industrialization of influence, where platforms and supporting institutions emerged to professionalize and globalize

the work of micro-celebrity. While most influencers still aspired to monetary success, an infrastructure of opportunity (and exploitation) waited to guide and enhance their celebrity strategies while extracting as much value from their labor as possible.

While much has been written about which specific tactics influencers use to cultivate relationships with their audiences, but perhaps none has been more thoroughly documented than portrayal of *authenticity* (Banet-Weiser, 2012; Marwick, 2013a; Marwick, 2013b; Duffy and Wissinger, 2017; Duffy, Miltner, and Wahlstedt, 2022; Hund, 2023). Authenticity is generally defined as a socially constructed quality that signifies "realness," vulnerability, and, sometimes, a sense of living an "ordinary" life. Audiences approve of authenticity in influencers, and are, for example, more likely to buy products from those they perceive as authentic (Pöyry et al., 2021). Influencers are compelled to manage perceptions of their authenticity among their audience, and one of the primary tasks of influencer labor is to construct the appearance of authenticity while often doing inauthentic things, like selling products one does not use and one's audience may not yet want (Abidin, 2017; Arriagada and Bishop, 2021; Driel and Dumitrica, 2021). Arnesson (2023) describes how audiences' evaluation of authenticity extends not only to commercial habits, but to influencers' assumed ideologies, and whether those ideologies accord with the products they are compelled to market. As I discuss further in the following section on politics and influencers, the drive to authenticity thus becomes a totalizing phenomenon, representing an certain audiences' ideology around the "normal life," and influencers' attempts to portray normal lives implicitly reify those ideologies as a condition of economic success.

Since the late 2010s, there has been an explosion of studies on figures termed influencers in most media studies journals concerned with social media.[3] In addition to greater interest and activities of product-based, commercial influencers, there has been what Bishop (2023) calls "influencer creep" into industries previously not thought to operate according to influencer logics. This has

[3]For example, as of April 20, 2024, *Social Media + Society* has published 124 articles with "influencer" in their text since 2016, none before; *New Media and Society* published 98 articles since 2018, none before; *Media, Culture, and Society* 32 since 2018, none before; *Information, Communication, and Society* 40 articles since 2016, one before.

been most thoroughly theorized in the realm of artists (Bishop, 2023), also called creative labourers (Duffy et al., 2021; Simpson and Semaan, 2023) and/or content creators (Ma and Kou, 2022; Ma, Gui, and Kou, 2023), whose creative economies have been subsumed into influencer-like incentives to self-brand and be "authentic." However, similar claims have been made (with some anxiety) about the practices of mainstream journalists, who are similarly pressured to adopt influencer practices of self-branding (Mellado and Hermida, 2021). As I describe in the following section, political pundits would soon find themselves subjected to the same system (Rothut et al., 2023). Effectively, influencers, marketing agencies, and platforms not only created the basis for an industrialization of their own industry, but a new economic system that could prey on precarious labourers across previously unaffected industries.

Understanding the industrialization of influencers is important to this dissertation because it informs the question of how to define and measure influencer-related phenomena. If influencers were only important inasmuch as they changed beliefs, then that would call for representation methods centered on belief change within their audiences. However, if we understand influencers as ordinary people motivated and sustained by economic incentives and external infrastructures, then this puts a renewed focus on their daily activities and strategies — even those strategies which do not seem to immediately change beliefs. Their activities, and the activities of marketers and media platforms which hoped to exploit them, create the necessary substrate for large-scale belief change. Media studies scholars thus show us that belief change on a mass scale is (perhaps unsurprisingly) wedded to the logic of markets and platforms.

2.2.4 Influencers, Politics, and the "Political" Influencer

The industrialization of the influencer thus bring us to the topic of this dissertation: the effects of influencers on political discourse. Emerging literature on this topic uses the term "political influencer" (Riedl, Lukito, and Woolley, 2023), although this can be misleading, as it suggests that there exist non-political influencers. All discourse has the potential to inform political values, and

indeed some types of influencer who may be popularly conceived as non- or apolitical engage in political discourse quite directly and frequently.

For example, as previously referenced, the choice to endorse a certain product often implies an endorsement of that product or company's implicit political stances and valences (Arnesson, 2023). Many have noted the porous relationship between beauty and wellness and the practice of far-right influencers, as narratives about the integrity of the body give way to conspiracies of outside forces attempting to degrade the (white) body (Baker, 2022; Moran, Grasso, and Koltai, 2022; Mastrangelo and Longo, 2024). Other scholars, particularly those studying Black bloggers, have noted how influencers have always blended the extraction of value from social media with self-conscious political activism (Steele (2021, chap. 4), Bailey (2021, chap. 3), Childs (2022). Abidin (2016) writes of how much-maligned influencer practices, such as the selfie photograph, can themselves be repurposed as subversive acts within exploitative platform economies, while Hearn and Banet-Weiser (2020) describe the implicit political valences of "glamour" culture itself among online influencers. Indeed, there is fair reason to believe that most political discourse online happens in not explicitly "political" contexts, given recent findings that most users follow few political influencers (as recently defined) on social media (Wojcieszak et al., 2022).

As Senft (2008, p. 117), Duffy (2017), Abidin (2016), Hund (2023, p. 32) and many others have noted, the female skew of micro-celebrity labor, at least in its early era, effectively *feminized* these types of political influencing, i.e. made them appear to some audiences as non-productive, voluntary, and concerned primarily with private, rather than public, life. We might add also that much influencer labor in U.S. political discourse is racialized and disregarded in the same way, given the paucity of non-English language studies of U.S.-based (political) influencers (Jaramillo-Dent, Contreras-Pulido, and Pérez-Rodríguez, 2022; Nguyễn et al., 2022). This is all to say that the theoretical lineage on "political influencers" which follows is the history of thought around a specific type of racialized, masculinized discourse, notable more for the self-conscious political aesthetic of its communication than for its dominion over political communication and political

belief. From early political bloggers to the contemporary far-right mega-influencer, white men as subject matter have dominated the research literature of the U.S. "political influencer," while only constituting a specific portion and type of political discourse online.

Arguably the direct antecedents of political influencers are political bloggers, and the greater idea of the blogosphere. Farrell and Drezner (2008) summarize the blogging era as one of the first instance of a truly networked political media, while lacking both the scale and profitability that we associate with influencer media in the present day. While the full spectrum of research on political blogs and bloggers is too vast to be summarized here, it bears mentioning as an intermediate development step towards the political influencer. The blogosphere era came to a sudden end[4] with the advent of Twitter and other large social media platforms (Bruns and Highfield, 2015), but the methods it pioneered, most prominently the mixed-method analysis of political communication over networked media (e.g. Herring et al., 2005), presaged future analyses of Twitter and other replacement platforms.

The creation of the "political influencer" as an academic term coincides with the dawn of the 2020s — previous uses were rare and colloquial, with "political" being used as a momentary adjective to the "influencer." Lewis' analysis of far-right YouTube influencers provides one of the earliest concrete definitions of a political influencer in 2018: someone who uses influencer strategies to maintain an audience and sell them an ideology, rather than a commodity (Lewis, 2018, p. 1). Three years later, Riedl et al. (2021) similarly conceptualized political influencers as influencers who may sell products, but also discuss "meaningful" topics, such as sustainability. Soon after that, a special issue on political influencers was published defining them expansively: expressing political speech full-time or part-time, profiting or not profiting from that speech, limited to platforms but possibly exceeding them, and encompassing of journalists, politicians, and "platform-built" influencers (Riedl, Lukito, and Woolley, 2023). An increase in explicit usage and citation of the term

[4]Using the search functions of the *ACM*, *Sage Journals*, and *Taylor & Francis Online* publication websites indicates a peak of the term "blogosphere" from 2009-2012, followed by a rapid decline.

"political influencer" in the past two years indicates a new concern with this archetype, contrasting with previous colloquial references to influencers who may or may not discuss politics.

As literature of political influencers is nascent, it is important to reaffirm that who gets to be a political influencer is contingent on their identity. First, to the extent that political influencing may not be as remunerative as, say, fashion influencing, and the extent to which all influencer labor is precarious and high-risk, those with more structural advantages are more likely to become successful (political) influencers (Bishop, 2021; Barbala, 2023; Christin and Lu, 2023). Second, whereas platforms generally attempt to support influencers in their design and moderation, they tend to disproportionately target users with marginalized identities, especially when they are engaging in political activism (Gray and Stein, 2021; Haimson et al., 2021; Duffy and Meisner, 2023; Abokhodair et al., 2024). Third, marginalized creators, especially discussing politics, are subject to increased harassment which likely limits their ability to maintain career success compared to dominant groups (Peterson-Salahuddin, 2022; Meisner, 2023; Han et al., 2023; Salehi et al., 2023). The primary takeaway for this dissertation is that methods which attempt to define political influencer status by their visibility on platforms likely underestimate the variety of true political influencer activity, particularly among marginalized communities.

2.3 Conclusion

These two theoretical lineages provide normative and practical frameworks for understanding who influencers are, what strategies they pursue, what incentives and constraints compel their behavior, and how their behaviors affect political discourse.

Threads in political communication provide a rich history of normative debates over the ideal functioning of media systems, and most importantly give us valuable context for understanding how people act in *groups* to produce and receive media. Decades of debate on opinion leadership and the two-step flow of communication provide historical context for contemporary influencers,

and provide normative stakes to the study of their effects. Research on publics supply a standard against which the group behavior of influencers can be measured, and provides a rich history of social media case studies against which this dissertation's studies can be compared. Scholars in audience theory provide important context for the primary method of this dissertation, the visualization of influencer communities through the mechanism of shared audiences, and highlight the deep well of audience activities which remain largely implicit in this method.

Media studies supplement systems-level research in political communication with ethnographic research on the actually-existing practices of influencers, viewed primarily through the prism of labor. It chronicles the emergence of a new form of mass communication; one which has its roots in pre-internet media, but which has radically transformed by the technical affordances and industrial exploitation of the internet. This research lineage provides insight into the economic conditions that compel influencers to behave as they do, and the strategies they pursue to maintain an "authentic," profitable brand. Recently, scholars have provided a tentative mapping of the *political* influencer, who is arguably the subject of this dissertation, although other strands of influencer ethnography highlight that the political effects of earlier, largely feminized influencers have likely been understated.

Chapter 3

Methodology

In this chapter, I make three types of contextualizing interventions for the studies that follow. First, I define how I think about the central objects of study in this dissertation, influencers and their communities. Second, I provide a review of the methods and data collection procedures I use across *all* of these studies. Individual studies have their own methods and background sections more specific to their contexts. Third, I interrogate ethical issues in this dissertation, as well as my own positionality. I begin, now, by discussing how I choose to conceptualize and represent influencers.

3.1 Definitions and Representations

3.1.1 Influencers

In the Background (Chapter 2), I reviewed three separate literatures which provided multiple, evolving perspectives on the contemporary figure known as the "influencer." In order to complete my own studies, I have needed to develop a working definition for influencers and their aggregates. The following section describes how I have come to define those definition, inspired both from the literature reviewed in the previous section, and empirical insights derived from my own case studies. We begin with the *influencer.*

An influencer is a person or entity that seeks to translate attention into different forms of capital within a highly metricized media environment.

Attention and Capital

I begin with the requirement that influencers be seeking attention on social platforms in order to exchange it for other forms of capital. Many influencers, clearly, pursue attention in exchange for financial capital, as more recent work in media studies has described (Duffy, 2017; Abidin, 2018; Hund, 2023). However, others more closely follow the opinion leader model, and pursue influencing to effect belief change or inspire action — political capital. Social media accounts for politicians and other political operatives are perhaps the clearest example of this motivation, using influencer tactics to promote ideologies, stigmatize marginalized groups, or win elections (Lewis, 2018; Marwick, 2021; Riedl, Lukito, and Woolley, 2023). Still further, some influencers seek attention simply for the affective capital of being popular and deriving joy from that popularity. For example, there is a rich scholarship describing the joy and community-building that online influencers in Black counterpublics facilitate (Lu and Steele, 2019; Steele, 2021; Musgrave, Cummings, and Schoenebeck, 2022), while others describe the sadistic joys of organizing misogynist hate campaigns (Cook, Schaafsma, and Antheunis, 2018; Just, 2019). Broadly, influencers frequently express enjoyment of positive interactions with their fans, despite other monetary incentives at play (Senft, 2008). The ultimate point being that attention and popularity are forms of capital which can be transformed into many other forms of capital, and so the pursuit of attention by an actor need not imply any specific financial motivation.

Metricized Media Environments

I then move on to the requirement that influencing happen in a highly metricized media environment. In the previous paragraph, I describe some quality which can be transformed into capital: attention. Attention is a type of interaction sought from audiences, and as described previously, media organizations, influencers, and ordinary users alike struggle to understand what attention

is, and which audiences it comes from. Marketing researchers have struggled for nearly a century to create metrics that construct and rationalize audience attention, working to increase the cost efficiency, sampling speed, and fidelity of their methods (Ang, 1991; Napoli, 2010).

The internet, and particularly the rise of large social media platforms and marketing tools, represented a significant change in how people were able to rationalize audience attention (Nelson and Webster, 2016). The new attention measuring systems that platform developers and marketers developed were inevitably flawed in countless ways and did not resolve questions around the *quality* of people's attention (Baym, 2013). Despite these flaws, they provided the ability to instantaneously monitor "attention" in a shared quantitative language of likes, shares, comments, etc. Media systems had become highly *metricized,* and this increased metricization allowed for new ways of interacting and extracting value from these systems.

I argue that the affordances of widely-accessible, mutually-intelligible, instantaneous "attention" measurement provided the necessary conditions for contemporary influencers' existence. With instantly-updated measurements, influencers could minutely tailor every aspect of their online existence to changes in their attention levels, resulting in the self-branding practices we see in contemporary influencers. Because these attention measuring systems are largely public and mutually intelligible, we see the conditions for an imagined community — and what I term below to be influencer communities — emerge around cultivating attention. Would-be influencers could observe the effectiveness of their peers' efforts at attracting attention, and generate shared pools of knowledge around how to become more effective influencers. The wide accessibility of the these measuring systems meant that many more people could become influencers than ever before, bolstering the development of this community of practice and its industrialization into an economic force that has now received widespread recognition.

Persons and Entities

Perhaps the most significant difference from contemporary definitions of influencer is my inclusion of non-human entities and organizations as potential influencers, as well as categories not typically considered influencers, such as politicians and journalists. I believe that these other categories, despite their unique qualities and histories, are increasingly subject to the logics of attention economies, and therefore can be productively analyzed jointly under the rubric of the influencer. Take, for example, news organizations. Recent research has emphasized how they increasingly prioritize their content via audience metrics on social platforms (Welbers et al., 2016), encourage their employees to use branding strategies typically associated with influencers (Holton and Molyneux, 2017), and engage interactively with audiences. Organizations, especially those created in the social platform era, increasingly talk on social media in the first-person, and interact in ways akin to human influencers (e.g. Deadspin, 2013; VersoBooks, 2023). Recent scholarly fascination with the "virtual," artificial influencer reflects that the principles of the attention economy need not be restricted to single persons representing themselves (Arsenyan and Mirowska, 2021; Conti, Gathani, and Tricomi, 2022). Rather than engage in boundary-work to understand what separates the influencer from entities which are subject to the same constraints and act in similar ways, I instead opt to consider them all under the influencer terminology for the purposes of this dissertation.

3.1.2 Influencer Communities

Having defined the influencer, I now define the term which is the focus of this dissertation: the influencer community.

Influencer communities are communities of influencers who, within certain discourse contexts, seek attention from the same audience.

Given the socially constructed (Ang, 1991; Napoli, 2010), imaginary (Litt, 2012) nature of audiences and attention, I understand community in this definition as an imagined community, per

Anderson (1983). Imagined communities are communities formed with people you may never directly interact with, but with whom you share a common identity that is assumed to develop from common experiences. Influencers are in imagined community with other influencers who, reading audiences like shared texts, are subject to the same attentional incentives. In some cases, influencers in these imagined communities can know and interact with one another on the basis of their shared labor under these audiences, creating formations akin to communities of practice (Eckert and McConnell-Ginet, 1992; Gannon and Prothero, 2018). But importantly, they need not interact with or even be aware of each other to be in a shared community. It is the fact that they labor under conditions of a shared attention geography, their audience, which brings them into these communities.

Influencer communities, per my definition, are contextual and non-exclusive. For example, in a political discourse environment, it follows that influencer communities are likely to be aligned across different shared partisan audiences. In a sports discourse environment, some of those same influencers may be aligned by their audiences' support for certain teams. It is plausible that one person or entity is an influencer to different audiences at different times, and thus has different memberships in different influencer communities depending on the topic being discussed. It should be noted that the same goes for audience members, who can support different influencer communities in different contexts: these communities are not segmentations, but fuzzy and overlapping partitions (Ang, 1991, ch. 4).

As I describe in the following "Methods Background" section (3.4) and especially in Study 1 (Chapter 4), I use methods from network science to create operationalizations of the influencer community concept. Shared audiences are represented by digital traces on social media platforms, where two or more audience members the same engagement with a target influencer. Influencer communities are seen as the iterative products of these overlaps between audience members, and represented as collections of users/nodes/influencers who all share the same audience, to some tolerable level of internal variation. In practice, I often filter influencer communities to the most popular

influencers, assuming that they are more effective representatives of the imagined community of influencers than marginal would-be influencers to whom relatively few audience members pay attention.

3.2 Datasets

Data for all four of these studies have been collected using the Twitter V1 API. This API can be used in different ways: in Studies 1 and 2, it is used to collect tweets based on keywords; in Study 3, it is used to collect the full timelines and all other user interactions with a set of tracked users; in Study 4, it is used to collect replies specifically to a set of tracked users. Study 2 has also used data collected with the Twitter V2 API, particularly for its ability to sample the full timelines of certain sets of tracked users. Study 4 also makes use of citation data from the Web of Science (Analytics, 2017) and the Internet Archive (Kahle, 1997). Studies 3 and 4 use hand-coded annotations of either tweets or URLs.

The social media platform Twitter has changed its name to X since the data for these studies were collected (Ivanova, 2023). This name change coincided with a variety of new Twitter policies and affordances which fundamentally change the nature of the platform and its data. Accordingly, I refer to this dissertation as a study of Twitter, rather than X. I refer to X when referencing the present-day incarnation of Twitter.

3.2.1 Data Limitations

My primary data limitations stem from the choices we make during data collection. Studies 1 and 2 use live term-based collections, but we cannot anticipate or even know every term which will be relevant to an event, and so we are always retrieving a partial sample biased by our pre-conceptions. Studies 3 and 4 use user-streams, which are similarly biased by which users we choose to follow – users discussing the 2020 election on the one hand, and elected officials on the

other. This focus on users involved in mainstream electoral politics is intentional, but necessarily hides from view users involved in non-electoral politics, and users who do not believe they are discussing "politics" as such. Our analysis is also limited to English-language communities, which excludes communities in the U.S. (particularly diasporic communities) who speak other languages (Nguyễn et al., 2022). We try to be explicit about the limitations of these data selection choices, but inevitably our limited perspective shield us from understanding the full consequences of these choices. We encourage dialogue on the creation of appropriate dataset selection to address meaningful social problems.

There is some ambiguity as to how much, precisely, the Twitter V1 API gathers (Tromble, Storz, and Stockmann, 2017; Campan et al., 2018). Our data collections are rate-limited, meaning that at times of high-activity, we are not collecting all tweets which match our query. This likely biases downward tweet counts during high-frequency events, and prior research has shown, concerningly, that tweets collected under rate-limit conditions are sampled in non-random ways (Tromble, Storz, and Stockmann, 2017). Furthermore, on term-based datasets, we collect data on many terms simultaneously for different studies, meaning that high-frequency events pertaining to an unrelated study could affect term collections in the present study. With this uncertainty in mind, I attempt to not make claims which depend too strongly on numerical counts, especially when based on small time periods or close comparisons.

As is discussed later in this chapter in the "The End of APIs," some of the data we make public can no longer be rehydrated, or can only be rehydrated at prohibitive expense. Additionally, as more time passes, more users will have either deleted, privacy-protected, or had suspended their accounts, intrinsically limiting the practical reproducibility even if X API access was to be returned at a future date. To mitigate this, we try to provide intermediate data products in additional to raw data, which enable some level of reproducibility to our work in the face of disappearing data.

3.2.2 On the End of APIs

Ostensibly, this dissertation is an attempt to lay out a set of methods for analyzing large datasets of social interaction data generated by digital media companies. Though it was not so often stated, this work and much of the literature it draws upon was enabled by the leniency of these companies; their willingness to share, or at least not prevent the collection of, their data about users and content on a massive scale. Recent developments, such as the shuttering of Meta's CrowdTangle data collection tool for Facebook and Instagram (Scire, 2024), the end of freely-accessible APIs for Twitter and Reddit (Ward, 2023), and the increased popularity of platforms which have never offered broad data access like TikTok and YouTube (Counts, 2023; Gottfried, 2024), make it appear as if this leniency has come to an end. Furthermore, industry-wide layoffs in "Trust and Safety" teams at social media companies, i.e. those most responsible for collaborating with academic researchers, indicate an ambivalent stance from these companies towards researchers (Goggin, 2024). In the particular case of Twitter/X, their leadership has explicitly retaliated against social media researchers with lawsuits, indicating outright antagonism (Bond, 2023). These new postures toward researchers provide little hope that data will become more accessible again in the near-future.

Others have noted this trend, dubbing our current research paradigm as the post-API age, or, more pessimistically, the data abyss and APIcalypse (Freelon, McIlwain, and Clark, 2018; Bruns, 2021; Vreese and Tromble, 2023). As a practical matter going forward, researchers have suggested leaning more heavily on non-API data scraping, compelling platforms and/or collaborating with them to release their data, and launching large-scale collaborations with users to access their data (Freelon, McIlwain, and Clark, 2018; Tromble, 2021). Some have noted that the post-API age is not necessarily unwelcome: that a glut of social media data arguably led to a decrease in standards for analyzing that data, and caused a paradigm where the availability of data drove mainstream research questions moreso than other, prosocial factors (Tromble, 2021). I agree on both points, in that I agree that scraping is an attractive method forward for post-API research, and that the

elimination of practical APIs will force researchers (including myself) to more closely pair the development of our research questions with our data collection plans.

However, the death of APIs does beget the question: how can we use the methods and insights of this dissertation, which were based in a data regime that no longer exists? I offer two answers. First, the methods described here, particularly co-engagement graphs, do not necessitate large datasets, and indeed Study 4 employs network projections on a small dataset of a few hundred manually annotated URLs. Second, I argue that the API era has allowed researchers like myself to offer, at their best, a view of unprecedented fidelity and breadth into large-scale sociopolitical systems mediated through the internet. This is particularly true for the visualization tools offered in these studies, which offer publicly-intelligible maps of information ecosystems which the public may never have access to again. By making the datasets and methods we have from this era public, we can offer a compelling comparative case study for future social media environments which will invariably have less data to work with. The data of this past era will guide us in understanding what is missing in future eras.

3.3 Ethical Concerns In Data Collection and Analysis

The use of large-scale interaction data collected without consent from private social media companies is problematic, even if technically admissible by institutional review boards (IRBs) in the United States. In theory, users read terms of service (ToS) documents before signing up for platforms which stipulate that their comments are public and could be used for research. In practice, most users historically would prefer their work not be used (Hudson and Bruckman, 2004; Fiesler and Proferes, 2018; Klassen and Fiesler, 2022), don't read or understand ToS (Fiesler, Lampe, and Bruckman, 2016; Obar and Oeldorf-Hirsch, 2018), or even on a basic level understand that their posts are public (Proferes, 2017). Researchers use anonymization to gloss over these concerns, but some subsets of users — particularly content creators, which are studied in this dissertation — hold a strong ethic of citation for their original creative works and prefer not to be anonymized

(Bruckman, 2002). Furthermore, if any user is directly quoted, even if they are lightly paraphrased, search engines make retrieving their original identity facile (Reagle, 2022; Mason and Singh, 2022). And if anonymization is broken and users are cited, they may receive outside harassment for their participation (Dym and Fiesler, 2020).

To put it mildly, there is no model of research which privileges active "participant" consent to which most contemporary large scale social media studies actually subscribe, or could even feasibly achieve at large data scales. With this in mind, I assess information disclosure on a case-by-case basis (per franzke et al., 2020), weighing the potential harms of privacy violations against the benefits of revealing this information within the context of each study.

This dissertation A) collects user-created information on social media, B) shares usernames in some studies, C) screenshots images from users' content, and D) quotes users' written words directly. Myself and my colleagues use several methods to mitigate privacy risks from these practices. First, with one exception, we only directly quoted, named, or shared screenshotted material from accounts who we considered public-facing: organizations, journalists or those whose self-publish e.g podcasts and other material, or politicians. The exception represented screenshotted content from an anonymous user who we did not name, but who had created popular anti-vaccine content. We consider for each user whether named or name-able users may face harassment due to our referencing their content — or whether providing additional publicity or authority to users who do harm may increase their public profile and social capital, thereby imperilling others (Arif, 2020, Chap. 3).

Our data representations include on some level the action of "ordinary" users who retweet these public-facing users, but we only publish aggregated data representations which severely obscure any individual action from those users. We never specifically describe the behaviors of individual non-public-facing users, and do not believe they can be identified from this data. In the interest of reproducibility, we do share large datasets of tweet and user IDs use to generate our data. This data cannot rehydrate data from users who have e.g. deleted or privacy-protected their accounts

or tweets.

Aside from individual methods, I have intentionally studied communities of popular users who make deliberate interventions into dominant public sphere to persuade mass publics in the United States. Given that the types of influencers I study, by every indication, intentionally seek widespread publicity for their online efforts, I view their expectations of privacy as lower than others. I deliberately do not collect interaction data from marginalized communities where I do not have membership or understanding, though data from these communities may be collected either by accident (because they e.g. use similar terminology to the influencers at hand) or when they intervene in dominant public discourses. I do intentionally reflect on the nature of studying dominant publics, rather than marginalized publics, and how these methods might travel differently when applied to different communities.

3.4 Methods Background

3.4.1 Network Science

The principle analytical tool in this dissertation is network science as applied to social networks (Barabási, 2013; González-Bailón and Lelkes, 2023). Network science is the practice of representing social phenomena as graphs: structured data in which nodes, usually representing entities, and are connected to other nodes via edges, usually representing actions or relationships. The field of network science is vast, with a robust literature detailing its methods in the abstract and in their application to a variety of fields including biology, computer engineering, and the social sciences.

It is the social sciences we are particularly concerned with in this dissertation. In these cases, nodes usually represent people, or entities they control like organizations or, in this case, social media accounts. Edges typically represent social relationships between people, and can be weighted by the intensity of the relationships they have formed. Nodes can also contain attributes irrespective

of their edges, such as gender in a network of people, and attributes derived from their edges, such as community membership defined by one's pattern of connection with others.

Social networks can exist in many different configurations and display a range of properties, but the internet-mediated social networks which we study have been shown to have specific properties relevant to this dissertation. One critical property is that of *preferential attachment*, wherein nodes with many connections are more likely to attract new connections (the "rich get richer" phenomenon, Jeong, Néda, and Barabási, 2003). Preferential attachment describes a basic model for understanding how networks can form with highly unequal connection hierarchies, i.e. with a few nodes hoarding almost all the connections. We can assume that a model with at least some element of preferential attachment forms the networks we see in this dissertation, where audiences of millions focus the great majority of their attention on mere hundreds of influencers. Each new member of a social media platform is far more likely to follow an already-popular influencer than they are to follow a random member of the crowd.

For the remainder of this section, I document the three principle methods from network science that I use in this dissertation: bipartite projections, community detection algorithms, and network visualizations. Each of these methods contributes in different ways to the core problem of analyzing the structure and behavior of communities of influencers. I also begin by describing how I represent influencer communities as networks in the first place.

Network Schematic

In each study, I initially consider Twitter data as a directed network. Directed networks have edges that are not reciprocal, with a sending node and a receiving node. In these case of Twitter data, User A interacts with User B through a retweet, quote tweet, or reply. Typically, user interaction data is aggregated over a time period, and the aggregations represented via edge weights. So, for example, if User A retweets User B five times, their directed edge has weight five.

In studies 3-4, I consider signed networks, where some edges have negative weights, indicating negative relationships, and some edges have positive weights, indicating positive or non-negative relationships (Tang et al., 2017). I consider retweets to be non-negative mechanisms of endorsement, following prior scholarship (Metaxas et al., 2015) and our own qualitative analysis. Quotes and replies are seen as interactions of mixed valence, with the potential to be negative.

Bipartite Projections and Coengagement Networks

In order to conduct analyses on these directed networks, I typically transform them into bipartite networks, and then bipartite projections. Bipartite networks are networks that contain nodes of two types. These two types make connections across types, but never within their own type. A typical case study of bipartite networks involves studies of fields in academic literature, in which there are author type nodes, there are literature type nodes, and the only connection allowed is when authors cite (form edges with) literature nodes (Newman, 2001). In the signed case, one might consider the act of customers reviewing products negatively and positively on a commerce website. Customers review products, but products do not review other products and customers do not review customers.

I consistently format the data in this dissertation in the bipartite mode. In my formulation, one node type is creators, and the other node type is the audience members that endorse or critique them. My construction goes one step further in imagining that the same person/account can be represented separately as both types depending on their contextual social role — in one situation, listening as an audience member, and in the other, speaking as a content creator. In other words, each user has one node for the active engagement they send out, and another parallel node representing the engagement they may receive from others.

A frequent desire in the analysis of bipartite networks is to analyze one node type in terms of the other node type. In citation networks, one may be interested in papers that are frequently cited together by the same scholars. The resulting clusters of papers could then represent notions of

academic fields. In my own analyses, I am interested in content creators who are consistently endorsed by similar audiences, which indicates a sort of shared theatre upon which their efforts to attract attention play out. To facilitate easier analysis and visualization of these networks, I create *bipartite projections* from my bipartite networks.

Bipartite projections create new graphs where one node type is implicitly represented in terms of the other node type. For example, if Influencer A and Influencer B are both retweeted by Audience Member C, then we can create a new network representation where Influencer A is directly connected to Influencer B, and the edge between them represents their shared retweet engagement from Audience Member C. Iterating this process over all possible Influencer-Audience interactions, one is left with a network of only Influencer nodes connected with edges which represent the shared engagement of different audiences members. Influencers with more audience members retweeting both of them have edges of higher weights, indicating a greater overlap of audiences. I call this graph a *coengagement network*, describe it fully in Study 1, and use it subsequently in Studies 2-4.

My motivation for using bipartite projections instead of the original directed networks is that I am interested in the effects of shared audiences on influencer behavior, and how audiences structure the shared conditions upon which influencers operate. A directed network could help us analyze many interesting questions about how and potentially why influencers interact with each other, but what drives my research questions here is not how influencers *perceive* each other, but rather how external audiences legitimize their behaviors, or critique and/or ignore those behaviors which they find distasteful. This is to say, in the sociological tradition (Goffman, 1959), I view influencer behavior as a performance, whose analysis necessarily includes analysis of the audience to which it is performed.

Clustering

The practice of community detection in networks, called more generally *partitioning* or *clustering*, is used in all four studies. This method classifies different groups of nodes into categories depending on their edge relationships with other nodes (Fortunato, 2010). Clustering is fundamentally arbitrary, because the definition of "related" among nodes is dependent upon the analyst's definition of relation, i.e. what constitutes a community (Javed et al., 2018). In this study, I use the modularity method extended to signed networks, which broadly means that communities have more positive ties within their community than without it, and more negative ties between other communities than within (Newman, 2006).

Clustering is a key practice because it allows me to group influencers into ideological communities based on their shared audiences, and then analyze behaviors typical to certain communities and interactions between different communities. I typically define a community on Twitter as a group of accounts who are linked together by the practice of the retweet (Metaxas et al., 2015). I view community assignment based on retweet interactions preferable to other methods, like assignment based on stated ideology or linguistic cues, for several reasons. First, influencers and indeed all people can not necessarily be trusted to state their ideological affiliations honestly, as it will sometimes be advantages to claim one affiliation while acting as another. Second, linguistic cues like word choice, hashtags, or expressed stances are sensitive to rapid changes in the meaning of words across datasets, as well as the deception mentioned in the previous section. Third, other metrics, such as persistent following relationships, can be deceptive, as users sometimes follow accounts that they would not endorse but still would like to monitor (e.g. ideologically-opposed news accounts and politicians). Clustering by retweet interactions essentially defers the problem of community assignment to the wisdom of the crowd, letting ordinary users essentially vote via retweet on which influencers are associated with which regardless of their stated views.

Because different clustering methods based on different definitions of node relation will output different groups, my results are somewhat contingent on my choice of algorithm. In my cases,

I use the Louvain algorithm, due to its availability in many software packages, its relative efficiency on large networks, its ability to automatically choose an optimal number of clusters, and its tendency to preserve large clusters instead of creating many micro-clusters, (Blondel et al., 2008). While there is no formal analysis of different clustering methods in these studies, I have not anecdotally found differences between different clustering algorithms which would affect downstream findings in these studies. I also choose analyses which would not be heavily affected by clustering differences, e.g. choosing to analyze the most central nodes which are almost always consistently clustered across algorithms.

Clustering is also usually computationally difficult, meaning that solutions tend to be approximate, and in some cases will differ between runs of the same algorithm on the same data with different parameter initialization. In some cases, we use boot-strapping methods to overcome these difficulties, essentially averaging runs of the same algorithm from different positions.

Visualization

A significant, if sometimes unstated, component of these studies is the practice of network visualization. Network visualization is the representation of graphs as visual node-and-link representations, where closeness between nodes tends to suggest similarity according to some metric (Brandes et al., 1999). This "suggests" phrasing is important, as many people mistakenly assume that the placement of nodes perfectly represents some qualities of those nodes, when in reality node placement is almost always a compromise between visual clarity and empirical validity (Beers, 2021). As such, network visualizations are not judged as data structures which could be used for e.g. statistical analysis, but rather by whether a visual analysis by their intended audiences could reveal valuable qualitative findings, deter misleading findings, and make explicit aspects of ambiguity which can provoke additional research questions. In this last sense in particular, network visualizations are often exploratory tools used to generate many hypotheses, rather than confirmatory tools used to explain the result of a single hypothesis test (Venturini, Jacomy, and Jensen,

2021; Jacomy, 2021).

All four studies have at least some network visualization method derived from the coengagement graph. In each case, I use the ForceAtlas2 force-directed algorithm implemented by Gephi, with a slightly varying set of parameters which tends to create visualizations with good properties for display in these studies (Bastian, Heymann, and Jacomy, 2009; Jacomy et al., 2014).[1] This particular method on generating visualizations was chosen for two reasons. The first is its availability in Gephi, a powerful, interactive, and high-resolution open-source visualization program that can handle a large amount of nodes and iterate quickly over different designs. An algorithm's implementation in an effective open-source software package is important not only for ease and speed of iteration, but also because it enables easy reproducibility of visualizations by subsequent researchers. Second, this way of producing visualizations tends to produce relatively dense, simple arrangements of nodes — largely blob-like shapes for clusters, with spans of different thickness between those clusters. Focusing on the representation of simple aggregate shapes highlights the community-detection network of these visualizations, which I find to be the most salient in my studies, and discourages the meticulous study of individual node-node interactions, which can be misleading in a force-directed layout. Further studies will be needed to evaluate the precise effectiveness and drawbacks of these visualization choices, although anecdotally, they have consistently provoked valuable exploration in both experts and lay-users alike.

I believe that network visualization has been an effective tool both for myself and for my audiences in exploratory analysis of study data and findings, but the practice should be approached with caution. Jacomy (2021) cautions against the practice of "storyletting," i.e. the practice of academic imagery unintentionally telling research stories that their authors did not intend, and also the "noema of big data:" the potentially misleading air of meaning that the appearance of network visualizations lends to potentially meaningless results. I have experienced both of these phenomena first-hand, and have published a workshop paper on both the popularity of my own network

[1] In the ForceAtlas2 program in Gephi, I use the "Stronger Gravity" feature, set edge weight strength to 0.4, and adjust the gravity and node scaling parameters to result in relatively square figures with few overlapping nodes.

visualizations on social media, and their subsequent misinterpretation (Beers, 2021). My work is made even more challenging by what Foucault Welles and Meirelles (2015) have described as the practice of audiences overlaying popular myths about their social worlds into social media visualizations of those worlds. I visualize what many people perceive (somewhat mistakenly) to be proxies for the behavior of people with certain political ideologies, and consequently their own myths about those ideologies, or the practices of politically-active people in general, are frames for which to view my data. I have not yet cracked the code for creating visualizations which anticipate both the inherent deceptive properties of network visualizations nor the effect of widely-held frames through which they are received, though each of my research efforts is an attempt to iterate on the failures of the previous.

Finally, we must remember that the act of create a network visualization of a community is an act of mapping, and act of mapping is an assertion of power (Crampton, 2001). Maps are not only filters through which a cartographer transforms and displays information, but artifacts which produce entirely new knowledge claims about what objects are worth mapping, and what those objects' boundaries are. For example, I purport to create visualizations of U.S. influencer communities focused on political discourse, but as I describe at many points, the choices and compromises I make in data collection and study design mean that I am certainly not creating visualizations of *all* influencer communities and *all* political discourse, if such a thing were possible. Rather, I visualize and tacitly reinforce the authority to speak of a certain sort of Twitter-mediated, typically elite, disproportionately white, and English-language discourse. While I attempt to discuss the implications of these choices in my studies, to use network visualizations is to take responsibility for the consequences of the knowledge claims that these visualizations make.

3.4.2 The Case Study

This dissertation is primarily a methodological intervention, but each of these chapters pursues these methodological interventions through the mode of case study. Crediting and adding to

Feagin, Orum, and Sjoberg (1991), I define case studies as in-depth, multi-faceted qualitative *and* quantitative investigations of a time-bounded social phenomena. Here, the phenomena is that of the influencer community, represented through network science, during the early 2020s of U.S. political discourse (mostly) on Twitter.

Feagin, Orum, and Sjoberg (1991) describe several benefits of the case study compared to narrower, hypothesis-driven work. It studies people in naturalistic settings, embracing the complications of phenomena which are attempted to be simplified by traditional experimental research questions. These complications can lead to more holistic descriptions and more productive extensions of theory than narrowly-delimited hypothesis tests acting across abstract or multiple contexts. The study of phenomena in time-bounded contexts also allows for the integration of historical analysis to the phenomena in question, which is particularly essential when studying political discourse.

I diverge from Feagin, Orum, and Sjoberg (1991) in that I adopt a primarily mixed-methods approach, integrating more quantitative than qualitative methods in my analysis of influencer communities. This approach is derived from my mentorship in the field of crisis informatics, where the practice of using "multi-dimensional signatures" and data science to perform case studies of individual crises (Maddock et al., 2015). I certainly stress the network analysis component of work with multi-dimensional signatures (compare to Arif, Stewart, and Starbird, 2018), which others have already noted are a natural fit to the case study format due to their relationality and complexity (Gummesson, 2007). For the remainder of this section, I document two other methodological inspirations I bring to the case study: trace ethnography, and historicism.

Trace Ethnography

My method for generating case studies of influencer communities is closely related to trace ethnography, the practice or attempt to use seemingly thin informational records to reconstruct thicker descriptions of online communities (Geiger and Ribes, 2011; see also "hashtag ethnography," Bonilla

and Rosa, 2015). In a broad sense, I deploy trace ethnography when I analyze large dataset of Twitter posts, impossible (and useless) to read individually, using ranked metrics and network visualizations. The goal of the metric is make salient certain social dynamics; to pick a needle of salience of a haystack of social information. The goal of the visualization is to map the boundaries and interchange points between communities, which are almost impossible to see from within a community. I argue that while trace ethnography is certainly not the only useful way to understand influencer communities, it provides certain perspectives which are difficult or impossible to recover from the inspection of thicker sources such as interview or textual data. My intention is to complement such analyses of thicker sources, rather than to use the privileged status of "big data" to overwrite them.

There must be a note of caution when describing this methodological inspirations. The use of the word ethnography can, to those familiar with more traditional forms of ethnographic research, imply a much greater engagement with the communities I study than I have so far displayed. Indeed, the most relevant community consulted in this work is simply the University of Washington and my colleagues within it — a limitation I document more thoroughly in the following Positionality section. When I propose methods in this dissertation, it is not a proposal that they should be used in isolation, but rather a hopeful invitation to recombine them with different methods that have different advantages, including greater community engagement. As Geiger (2011) references in their foundational work on trace ethnography, one of the most effective uses of trace ethnography is suggesting to practitioners of thicker, more-embedded ethnography where to look next.

Description and Historicism

My approach in study design is inspired by particularly *historical* case studies of technologies and the communities of practice which arise around them (Soden et al., 2021). These studies are not only case studies of certain types of social formations, patterns of technology use, or dynamics

of information transmission. There are studies of these things specifically within the historical context of the U.S. in the 2020s, and particularly within the communities of scholars attempting to make sense of them. Outcomes and dynamics in these case studies are emergent from historical conditions, rather than invariant laws, without being precisely determined by them. Taking an approach inspired by historical work has affected my methods in this dissertation in two ways: my privileging of description and representation over explanation and hypothesis, and my corresponding attention to the narrative history of theory development in my literature reviews.

First, there is a significant amount of exploration and description in these studies. Rather than denigrating the practice of "mere description," I follow in the tradition of reclaiming descriptive, narrative, and analytic work, as opposed to problem-driven or explicitly explanatory work (Megill, 1989; Gerring, 2012). The process of mapping publicly-described phenomena to useful representations in scholarship is at least as valuable as what comes after, the testing of propositions within those representations. In this case, I claim significant value in the process of findings ways to represent influencer communities, and the myriad ways they interact with each other and change over time. Within the theoretical framework of description proposed by Gerring (2012), my work highlights *typologies* of influencer communities, *associations* between these communities, and *particularities* within these communities in the form of qualitative inspection of quotations.

I have described in the prior section how network visualization can be a productive descriptive tool, not only for illustrating the results of hypotheses, but for provoking new hypotheses through the generation of new data representations. More generally, in each of these studies, I provide a model of influencers and their behaviors derived from network science, and then use a variety of methods to assess different *types* of explanatory questions which may be asked of such a representation. The goal here is to describe; to lead by example in the practice of asking questions of these representations; and to start the process of creating, as historiographer Megill (1989, p. 646) describes, "a vast recounting, into which explanations are stuck like pins into a pin cushion."

Second, in light of my descriptive focus, I attempt to understand the history of methods development not as a series of attempts to improve a metric within a given task, but rather an ongoing, collaborative effort among scholars to derive *representations* which facilitate new questions and, generally, valued knowledge production. The question of representation and description, as Gerring and Megill note, is often a question of the normative values within a profession: what is worth representing, and what is not. Accordingly, my literature reviews focus on the longitudinal development of different representations, and what "problems," i.e. norm violations, they attempt to address in their respective literatures.

The Background chapter of this dissertation is my most direct attempt at such a theoretical narrative. It presents not a typical methodological review of attempts to solve a problem in measurement, but a narrative lineage of a concept-to-be-measured from its prior forms. The prior concepts of radio demagogues, opinion leaders, and micro-celebrities each represent different communities of scholars attempting to grapple with something that would soon converge and evolve into the contemporary influencer, and to understand their struggles is to understand the primary tensions to which the measurement methods I propose must attend. A similar method is followed particularly in Studies 2-4, where the concepts of community change, negativity, and citation are viewed as historically contingent concepts, rather than stable entities one attempts to measure against.

3.5 Positionality

Within this chapter and in the forthcoming studies, I have tried to faithfully describe my methodologies, and when possible follow principles of open science by making code, data, and research questions publicly available for external audit. However, only describing one's explicit choices in study design can run the risk of masking just how contingent one's research results are on their personal position, with respect to their studies' participants, chosen theories, and broader social-academic context (Hampton, Reeping, and Ozkan, 2021; Liang, Munson, and Kientz, 2021).

Accordingly, I report some of my own background here and reflect on how my positionality may affect my choices in the studies of this dissertation. I view the practice of reflecting and reporting on one's positionality as a method coequal with the prior methods of the network science and the case study. Knowledge is inseparable from the context in which it was created, and to evaluate a contribution to knowledge is to evaluate both its external contexts (positionality) and its internal logics (network science, the case study, etc.). Much of this positionality statement is inspired by feminist and particularly black feminist theorists of standpoint, who hold that knowledge production is inseparable form one social position and intersecting identities, and indeed that marginalized identities are a privileged position from which to create knowledge about marginalized groups (Haraway, 1988; hooks, 1989; Collins, 1989; Collins, 1990; Harding, 1991). Insights from standpoint theory here are slightly inverted, however, as I describe how aspects of my standpoint in many ways conform to what Collins (1989, p. 751) calls the Eurocentric masculunist approach of knowledge validation. Accordingly, this positionality section reflects on ways that being embedded in such a knowledge validation environment has affected this work, and should affect is interpretation.

3.5.1 Machine Learning, Big Data, and Industry

I came to work at the University of Washington's Human Centered Design and Engineering Department by way of the Quantitative Tumor Imaging Lab at the Massachusetts General Hospital, where I worked in machine learning research relating to medical imaging. Historically, this work came at a time in which interest and funding for machine learning research was rapidly growing in both industry and academia, and indeed that growth has accelerated to a fever pitch at the moment of writing (Vynck, 2024). My early experience with quantitative and particularly machine learning methods, coupled with this widespread institutional interest in these methods, creates natural publication and funding incentives to apply such methods regardless of their suitability to research priorities (Messeri and Crockett, 2024).

As many have observed, the growth of machine learning as a field is caused at least as much by industry power consolidation as it is by any actual societal benefits or technological advancements (Whittaker, 2021). Representatives of private industry themselves frequently participate in academic research in big data and machine learning, arguably performing an agenda-setting strategy to warp academic theories' to their funders' needs (Young, Katell, and Krafft, 2022; Abdalla et al., 2023; Gillespie, 2023; Burrell and Metcalf, 2024) It follows that my work, particularly Study 3, is necessarily a result of an industrial project designed to realize machine learning as a method, though my research has taken no direct funding from industry sources.

3.5.2 R1 Research and "Public" Scholarship

Similarly, I have completed my PhD at an R1 research institution at a well-funded lab which has received considerable media attention (e.g. Leingang, 2024) and therefore prestige within its academic ecosystem. Working at an elite lab within an elite institution — "elite" used here as a constructed category, rather than an evaluation of quality (Liu, 2011) — has several systemic effects on what research one is incentivized to create.

The first is that the diversity of perspectives is reduced at elite institutions, as only those with some level of privilege are allowed to access their resources. This creates a pressure to narrow the focus of one's work to elite interests which will be appreciated by your colleagues, your institution, your funding sources, and your larger ecosystem of similar R1 universities. More generally, it creates pressure to adapt research questions and contexts to ideologies of the academic elite.

Precisely determining the contours of the ideologies which shape sociological research at elite institutions is a larger project than I can commit to here. However, I have observed that our work has frequently been driven by the language of "crisis," due not only to its history in crisis informatics, but also due to the fascination that crises generate in dominant publics. Following Strolovitch's aptly-titled *When Bad Things Happen to Privileged People: Race, Gender, and What Makes a Crisis in America* (2023), I understand crises as socially constructed, and in popular discourse

crises tend to be constructed according to the interests of the powerful. Per Strolovitch (2023, p. 199), "...dominant understandings and assertions about when we have entered or exited a crisis are shaped as much by conventions about what is normal and whose pain is tolerable as they are by the severity of the problems at hand."

More specifically, I have observed that certain conceptual frameworks intertwined with narratives of crises, such as misinformation, polarization, and hate speech have been particularly consonant with the priorities of this elite ecosystem. Each of these concepts has been critiqued for servicing a worldview and neoliberal ideological project typically held by elites (misinformation, (Mejia, Beckermann, and Sullivan, 2018); polarization, (Kreiss and McGregor, 2024); hate speech (Spade and Willse, 2000; Conrad, 2014)). Accordingly, to the extent these studies encourage or adopt these frameworks, their acceptance for publication and citation by fellow elites is likely increased.

3.5.3 Elite Membership

Finally, I share many of the qualities of those disproportionately represented among those who receive PhD dissertations in my fields, and/or are disproportionately promoted within those fields (Morgan et al., 2022). I am white, a second-generation college and graduate degree student with an older sibling who pursued a PhD before me, and from a relatively wealthy background, to reference just a few examples of the structural advantages which have inevitably shaped the way I have designed this dissertation's studies. Particularly, exposure to elite political discourse and media figures in my upbringing, not only through media but increasingly through my peers and colleagues, naturally predisposes me to find the activities of these media elites more compelling in my research.

Many have documented the extent to which social media populations are not reflective of offline populations, usually favoring the perspectives of the privileged (Wojcik and Hughes, 2019; Hargittai, 2020; Shugars et al., 2021) (though not always — see critiques of the "digital divide" hypothesis, Tynes and Mitchell, 2014; Hamilton, 2020). Indeed, the set of influencers I study is

disproportionately white, especially at the highest levels of popularity. Similarly, Collins (1989, p. 753) has described how those with dominant standpoints, particularly white and male, generally control the knowledge validation process, promoting certain academic products that align with their "common sense" and rejecting others which stem from their lack of experience with or as part of marginalized communities. It stands to reason that research focused on mostly white influencers may appeal to a discipline in which knowledge validation procedures are still controlled in the majority by white academics.

There is, of course, still merit to studying the machinery of white U.S. internet media at the highest levels, as most people unwillingly live in the world that whiteness creates. My personal motivation for studying these topics is to expose and hopefully undermine aspects of this largely white media ecosystem. But I caution readers against losing sight of that history is as much written by those at the margins as it is by those at the center, and that a preoccupation with the center can sometimes simply reflect the subject position of the majority of tenured academics who occupy it.

3.5.4 The Dissertation as Branded Content

I describe my positionality not as an apology or a spontaneous unburdening of guilt, but to emphasize again: knowledge is inseparable from the context in which it was created. While I have raised concerns stemming from aspects of my positionality, I have of course attempted to work against these problems in the process of my writing, and have learned upon prior mistakes over the course of my dissertation. That I have done so does not obviate the reality of this work, which is that it stems from an academic environment which privileges quantitative, industrial, neoliberal, white, and elite perspectives. The dissertation as artifact carries the history of its own production with it, and to better understand the circumstances of its production is to better evaluate the validity of its claims.

Many who study influencers have noted the irony that academics themselves engage in influencer practices (Duffy, 2017, p. 230). Academics engage in branding, and universities often provide guidance in the finer points of online influencing (Pavlenko, 2018; *Lose your Fear of Academic Twitter* 2021). It follows that academics' publications are branded content, and that we too are producers of media. In the spirit of reflexivity then, we should always attend to the circumstances in which our media is created, including the matrices of power and domination which inevitably precondition its generation and reception. To quote Hall (1982) on the politics of television broadcasters:

> ...in the critical paradigm, ideology is a function of the discourse and of the logic of social processes, rather than an intention of the agent. The broadcaster's consciousness of what he is doing—how he explains to himself his practice, how he accounts for the connection between his 'free' actions and the systematic inferential inclination of what he produces—is indeed, an interesting and important question. But it does not substantially affect the theoretical issue. The ideology has 'worked' in such a case because the discourse has spoken itself through him/her. Unwittingly, unconsciously, the broadcaster has served as a support for the reproduction of a dominant ideological discursive field ("*From the 'reflection of consensus' to the 'production of consent'*", paragraph 6).

Chapter 4

Study 1: Coengagement Networks for the 2020 U.S. Election

Contributing Authors: Anna Beers[a], Joseph S. Schafer[a], Ian Kennedy[b], Morgan Wack[c], Emma S. Spiro[d], Kate Starbird[a]

[a]University of Washington, Human Centered Design and Engineering

[b]University of Illinois Chicago, Sociology

[c]Media Forensics Hub, Clemson University

[d]University of Washington, Information School

4.1 Preamble

This chapter was originally published in Beers et al. (2023a), and the text and figures outside of this preamble section are unmodified. This preamble serves as additional contextualization for this study within the context of this dissertation, and personal reflection on my part for where this work has since travelled. I use single-spacing in this preamble to offset it from the main text.

This study was originally inspired by work done for the University of Washington's Election Integrity Partnership, and multi-institution collaboration which included my lab, the Center for an Informed Public (Center for an Informed Public et al., 2021). It was specifically derived for visualization work I provided for Tollefson (2021), which features similar data and visualization styles.

My work for Tollefson (2021) is particularly relevant for this dissertation because it formed the basis for influencer panel I would later track for Study 3.

The context of this study within the dissertation is perhaps the most clear: it proposes a method for visualizing the structure of influencer communities, and an implicit definition of them, although neither is thoroughly elaborated with this study. I came to the understanding that my topic of study *was* influencers partly through the work of this study (and prior research efforts in the Election Integrity Partnership), which revealed to me that Twitter audience attention were mostly focused on platform-first influencers as opposed to mainstream sources and politicians. Subsequent studies would see me adapting my frameworks to include influencers in different ways, with an explicit focus on influencers in the chronologically latest study, Study 3.

Perhaps the second biggest impact this study had on future dissertation studies was the delineation of influencer communities which would prove stable both in the following studies and in unpublished, exploratory analyses. The pro-socialist community identified here, for example, also appeared in a tentative analysis of Russia-Ukraine war discourse on Twitter, with provocative connections to pro-Trump influencer communities. I am currently compelled by the extent to which the three-community structure is a generally latent structure of U.S. political political discourse, or whether by contrast it collapses in e.g. youth media, or non-explicitly "political" discourse.

Per my comments in Chapter 3, this mapping project has some notable exclusions. Unpublished work derived from this study shows how, with certain filtering methods, the influencer communities visualized here can be decomposed into different constituent communities. These include (part of) Black Twitter and news publics within the pro-Biden cluster, and white nationalist and evangelical communities within the pro-Trump cluster. It is an interesting area of future work for me to create flexible coengagement representations that could forefront influencer sub-communities. Work on network sparsification and network "backbones" is among the most relevant in the literature here (Nocaj, Ortmann, and Brandes, 2015; Nick et al., 2013).

If there were any findings from this study that went unexplored, I would say it is the representation of what I called, perhaps misleadingly given the overlap with Squires (2002), satellite communities. Influencers who translate information from dominant U.S. "political" publics, such as those from other countries literally translating, or politicized members of not-traditionally-political publics, like pop music, could be valuable ways to study the downstream effects of this dominant public. Given that most people do not directly follow political elites (Wojcieszak et al., 2022), they may still be feeling their effects through the translation efforts of these key intermediary influencers.

4.2 Abstract

The 2020 United States (US) presidential election was — and has continued to be — the focus of pervasive and persistent mis- and disinformation spreading through our media ecosystems, including social media. This event has driven the collection and analysis of large, directed social network datasets, but such datasets can resist intuitive understanding. In such large datasets, the overwhelming number of nodes and edges present in typical representations create visual artifacts, such as densely overlapping edges and tightly-packed formations of low-degree nodes, which obscure many features of more practical interest. We apply a method, coengagement transformations, to convert such networks of social data into tractable images. Intuitively, this approach allows for parameterized network visualizations that make shared audiences of engaged viewers salient to viewers. Using the interpretative capabilities of this method, we perform an extensive case study of the 2020 United States presidential election on Twitter, contributing an empirical analysis of coengagement. By creating and contrasting different networks at different parameter sets, we define and characterize several structures in this discourse network, including bridging accounts, satellite audiences, and followback communities. We discuss the importance and implications of these empirical network features in this context. In addition, we release open-source code for creating coengagement networks from Twitter and other structured interaction data.

4.3 Introduction

The 2020 United States (US) presidential election was — and has continued to be — the focus of pervasive and persistent mis- and disinformation spreading through our media ecosystems, including social media (Benkler et al., 2020; Center for an Informed Public et al., 2021). Efforts to understand these dynamics have driven the collection and curation of large social media datasets, and the subsequent production of large, directed network representations of social interactions to make sense of them (Abilov et al., 2021; Kennedy et al., 2022). But such network datasets can resist intuitive understanding. In large network datasets, the overwhelming number of nodes and

edges present in typical representations create visual artifacts, such as densely overlapping edges and tightly-packed formations of low-degree nodes, which obscure many features of more practical interest (Schulz and Hurter, 2013; Nocaj, Ortmann, and Brandes, 2015) (Figure 4.1). In the case of the US presidential election, one feature of particular interest is the functional level of interaction between different political communities who, due partly to pervasive misinformation spread in this country's right-wing media ecosystems, no longer share a common understanding of the election's outcome (Pennycook and Rand, 2021; Reuters, 2021). Critical to understanding these inter-community interactions is characterizing the role of platform elites, who are responsible for a disproportionate share of election-related misinformation (Center for an Informed Public et al., 2021).

Here, we present an extensive case study on a dataset of English-language Twitter posts relating to the 2020 US presidential election. We take advantage of the interpretative capabilities of coengagement networks, which are similar to the co-citation networks widely used in bibliographic scholarship. This dataset, totaling 585M retweets collected from September 1st, 2020 to December 18, 2020, contains tweets referencing generic English-language terms related to voting and the election, with a focus on tweets relating to election misinformation. In practice, this dataset contains public discourse related not only to the presidential election, but also discourse related to the persistent and false claims that the results of the election were fraudulent. We create and interrogate three different coengagement networks of retweets filtered under different parameter sets, describing via a mixed-methods analysis how the salient features of these networks correspond to different discourse phenomena. These phenomena include bridge nodes, users that are retweeted by multiple and disparate audiences; satellite audiences, groups of detached users which connect to mainstream conversations in very specific ways; and followback clusters, unique and highly active groups of users that incessantly retweet each other and very specific mainstream accounts. Our analysis of followback clusters particularly shows how Twitter's much-noted mass account removals in the wake of the 2021 attack on the US Capitol Building particularly affected these

followback groups.

Our empirical and methodological contributions together are themselves a case study in the proposed *triangulation* analysis method, where the intersection of understandings from multiple, sometimes contradictory network generates greater knowledge than any one network alone (Doreian, 1988; Brandes, 1999). We conclude this paper by discussing the advantages of coengagement networks over other social network formats, the importance of triangulation as a method for analysis of social networks, and future extensions and ethical considerations for using such a method.

4.4 Background

4.4.1 Mis/Disinformation, Platform Elites, and US Presidential Elections

Researchers have demonstrated the critical value of networks and network visualizations in efforts to identify key actors in political mis- and disinformation campaigns (Starbird, 2017; Starbird, Arif, and Wilson, 2019). Even within work in this domain, however, the terms mis- and disinformation themselves have been variously defined (Jack, 2017), sometimes eschewed in favor of the broader term "influence operations" (Wanless and Pamment, 2019), and sometimes even criticized as a contemporary moral panic (Mejia, Beckermann, and Sullivan, 2018; Carlson, 2020). For this paper, we simply define mis- and disinformation of interest as the unintentional and intentional spread of false or misleading claims that the results of the 2020 US presidential election were fraudulent. While previous research on disinformation in the 2016 US presidential election focused on foreign interference (Lukito, 2020), recent analyses of disinformation in the 2020 US presidential election have focused on domestic right-wing campaigns coordinated by elites on social media and beyond (Benkler et al., 2020; Center for an Informed Public et al., 2021). Recent research has shown how platform elites vary between different political groupings in the US, and have highlighted their role in spreading misinformation during the COVID-19 pandemic (Gallagher et al., 2021).

4.4.2 Challenges in Visualizing Social Networks

Network visualizations of large social data can provide valuable insight into the structure of online conversations, and these visualizations have become increasingly popular as representations of computational social science's promise (Foucault Welles and Meirelles, 2015). A common goal in social network visualization is to highlight influential nodes and characterize the relationships they hold with one another (Freelon, McIlwain, and Clark, 2016; Arif, Stewart, and Starbird, 2018; Stewart and Spiro, 2021). The simplest approach is to visualize the network in its observable entirety, with nodes representing user accounts and edges representing interactions between accounts. However, the large size of social media datasets, now often numbering in the millions or billions of nodes and edges so defined, can be intractable to visualize and render the exercise of doing so meaningless. Large social networks often encode multiple and seemingly contradictory dynamics at different scales, further exacerbating the difficulty in faithfully representing these phenomena to scholarly peers and the lay public (Jacomy, 2021). These representational ambiguities can become especially misleading in the case of social data, where the lay-public often has strong priors about what to expect from the social world (Foucault Welles and Meirelles, 2015).

A common solution to this problem of large graphs is to heuristically filter unimportant nodes (e.g. with low node degree) or edges (e.g. with low edge weight) until the visualization reaches a tractable size (Ham and Wattenberg, 2008; Dianati, 2016). Identifying those unimportant nodes and edges is a significant challenge, as the concept of *importance* is highly contingent on the interpretative aims of the researcher, and individually unimportant nodes may yet in the aggregate encode relevant structural information. Furthermore, there may be no single definition of node importance which addresses the full spectrum of phenomena represented by a social network. One innovation of particular relevance is the co-citation network developed in bibliographic network science, in which two published articles or authors are connected in a new work if other articles cite both of them together (Small, 1973). Importance filtering based on frequency of interaction in co-citation networks is frequently implemented, and its effects on apparent resulting

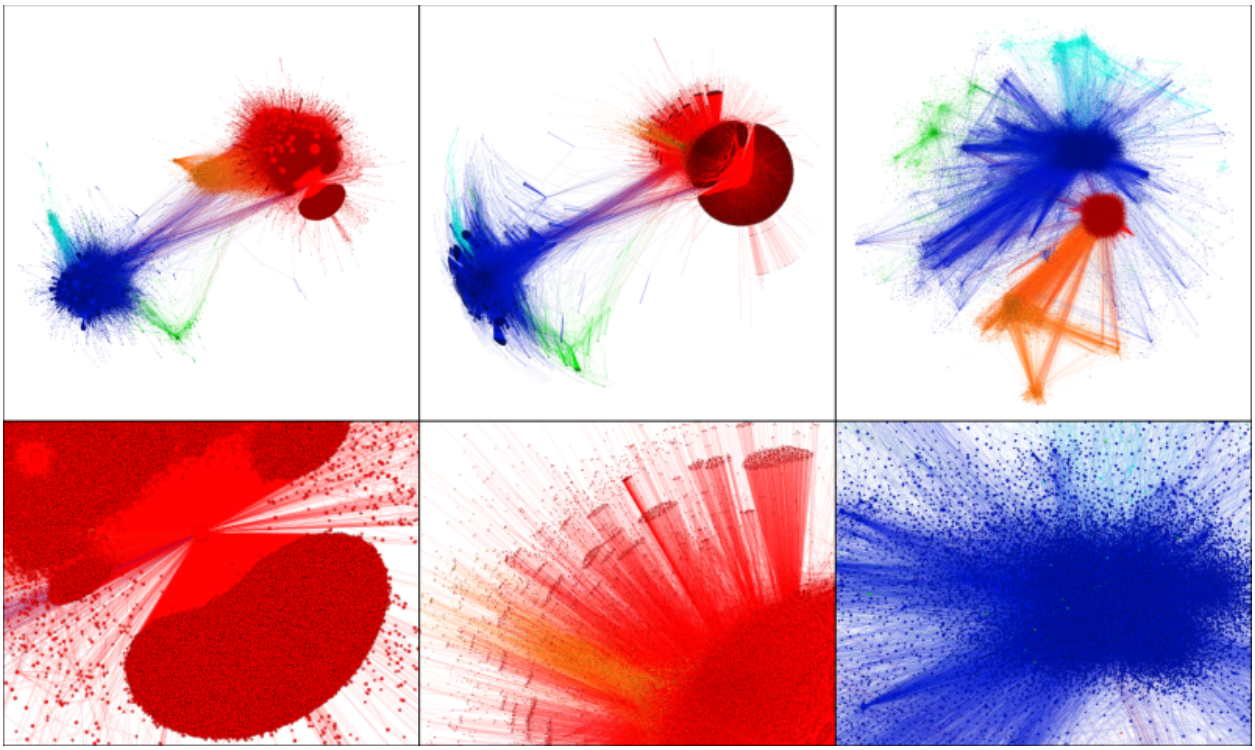

FIGURE 4.1: Directed retweet networks with typical node layout algorithms. Here, each node is an account, an edge represents a directed retweet, and edge weights represent the number of retweets. Edges with fewer than 50 retweets have been filtered out to aid visualization. The following layout algorithms as implemented in the software package Gephi are used from left to right: ForceAtlas2, YifanHu, OpenOrd. The bottom row contains close-ups of the same figures in the top row, highlighting dense node formations which obscure cluster interpretability. Dataset derived from Twitter data on the 2020 presidential election (size = 142K nodes, 424K edges).

clusters have been analyzed in (Shaw Jr, 1985). Co-citation networks have typically focused on scientific literature, although others have used similar principles in relation to the hyperlink structure of the web, most notably with Kleinberg's concept of online *hubs* and *authorities* in search engine retrieval (Kleinberg, 1999).

Here, we extend the concept of co-citation networks to social interaction data in what we call the coengagement network. Like co-citation networks, this method significantly reduces the number of nodes visualized in large datasets, while encoding information from missing nodes in visible

edges that can reveal significant relationships. Additionally, this method is also tunable, meaning that researchers can produce different visualizations according to different notions of node importance as defined by two interpretable parameters. A primary contribution is the application of co-citation principles to social data representing users, rather than documents, interacting with one another online.

4.5 Coengagement Networks

Coengagement networks are closely related to the method of projection in bipartite graphs. In a typical bipartite projection, a network with two types of nodes and no within-group connections (such as a network composed of researchers and the papers that they author Newman, 2001) is projected into one primary node type, with the remaining node type being collapsed into edge representations. This operation can productively reduce the number of nodes and edges to be visualized and analyzed, and preserves information about a primary node type while still retaining structural information from the projected node type. While we would not expect engagement networks among social media users to be bipartite, we aim to take advantage of the visual and analytical properties of bipartite projections, and therefore propose a transformation of non-bipartite graphs into bipartite graphs via the duplication of each node into two types: engaging and receiving nodes. We then project the resulting bipartite network such that engaging nodes are collapsed into edges between receiving nodes. In the context of users on Twitter, for example, this method privileges users with large audiences engaging in retweeting, commenting, liking, or even viewing, and defines relationships between users in terms of sharing similarly engaged audiences.

Formally, we define a directed graph $G = (V, E, w)$ with vertices V, edges E, and edge weights w. In the current case study, we interpret G as a collection of Twitter users (V) retweeting other users (E), with edge weights w defined as the total number of retweets from one user to another. We then define a new graph G' with vertices V' , which contains duplicate sets of vertices V_S and V_R that send and receive retweets, respectively. We similarly define a new set of undirected, weighted

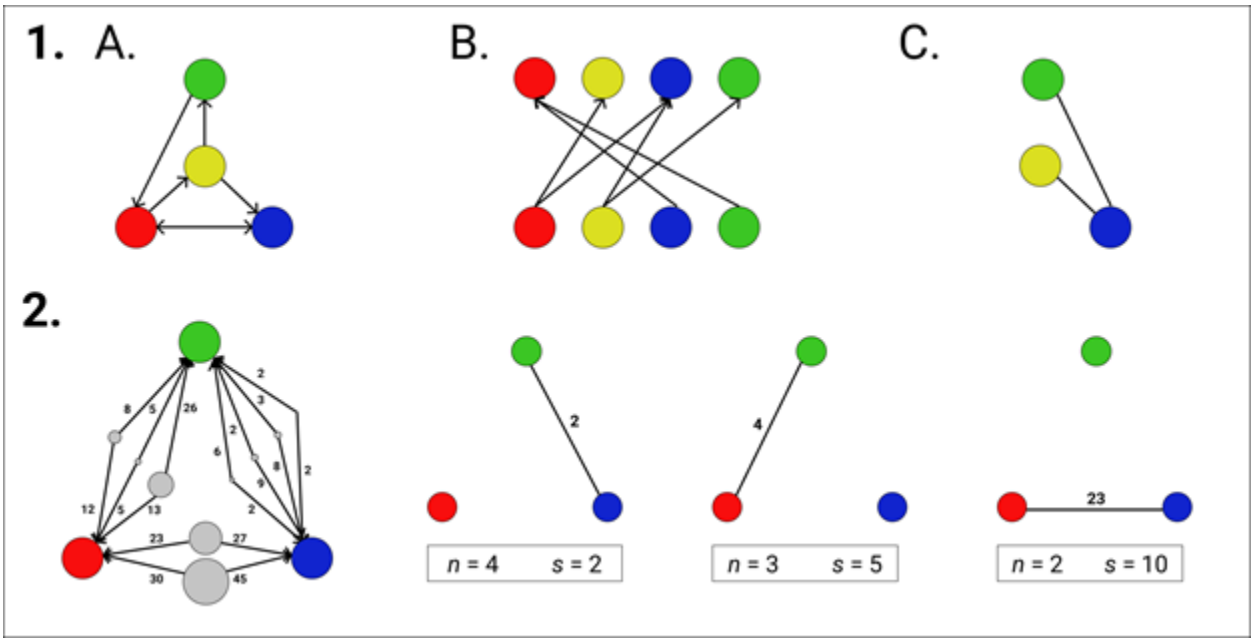

FIGURE 4.2: Schematic representation of coengagement visualizations. Panel 1) of this figure shows a schematic for transforming A) a unipartite directed graph into B) a bipartite directed graph, and then into C) a coengagement network. The blue node is linked to the yellow node because of their shared engaging node in red, and the blue node is linked to the green node because of their shared engaging node in yellow. Panel 2) shows the effect of the node filtering parameters on an example graph. Engaging nodes, colored grey, are sized by their average out-degree. Under different combinations of the filtering parameters n and s, different edges will result.

edges E' in G', where directed edges from V_i to V_j are represented as undirected edges from V_{Si} to V_{Rj}, with the same weights w'. In effect, each original user has a vertex representing the instances in which they retweet others, and a separate vertex representing when they are retweeted. We then define a projection of G' as X, such that vertices in X are the interaction-receiving vertices V_R, and edges in X are defined such that the edge weight between any two vertices i, j in X is the number of vertices in V_S that have defined edges to both V_{Ri} and V_{Rj}. This final vertex set in the projected graph represents users that are retweeted and draws edges between them when they are jointly retweeted by at least one other user. We finally define two edge filtering parameters n and s on the resulting graph X. Specifically, an edge between two users in X is defined if at least n other users have retweeted both users at least s times each. The parameter n represents a minimum

diversity of users retweeting two users, while s represents the minimum volume of retweeting a user must do to be considered in n.

The n and s parameters, which in practice control the number and distribution of edges in a co-engagement network, are a powerful tool for targeting specific visualizations. These edge filtering parameters allow researchers to shape the output of their networks along two important, yet distinct, qualitative dimensions by modifying the distribution of edges between users. When researchers filter with a higher n value, influential users are related only if they attract engagement from large, diverse audiences, a typical goal in influential user analysis. When researchers filter with a higher s value, nodes are instead related by audiences that frequently retweet their content with a dataset, which can reveal dedicated rather than transient audiences. Different ratios of n to s can reveal other relationships: low n with high s can make highly active and coordinated audiences more salient, whereas low s with high n make more salient those infrequent instances in which content is shared widely across different communities.

In what follows, we present a series of case studies using a dataset of tweets related to the 2020 US presidential election. Specifically, we use a dataset of 585M Twitter posts containing substrings related to voting and the election ('vote', 'voting', 'mail', 'ballot', 'poll', and 'election.'), define nodes as the authors of these posts, and (unprojected) edges as retweets from one user to another. We do not include quote tweets. We show that depending on the choice of the parameters n and s, different clusters of influential nodes can be distinguished, and different forms of qualitative analysis can be applied. In doing so, we demonstrate the practical implications of coengagment networks and their interpretation. Empirically, these case studies offer a unique look at engagement during the 2020 presidential election.

In the first case study, we choose a very high value of n and $s = 1$ to create a network that shows a broadly two-part structure to Twitter discussions around the presidential election, aligned with

pro-Trump and pro-Biden accounts. The low s parameter highlights transient instances of high-volume crossover between these two groups but does not necessarily represent *sustained* engagement across these groups. We then choose a parameter set with much lower n and slightly higher s, to illustrate how a third pro-socialist grouping becomes salient at different audience sizes, and how some crossover nodes do not hold sustained engagement. We end with a third case study at very high s and very low n, to highlight two new followback communities that become salient when active, sustained engagement is prioritized over large audiences.

4.6 Findings

4.6.1 Case 1: Bridging Between Clusters at High Audience Sizes

In Case 1, we generate a coengagement network where node relationships are defined by low restrictions on retweet frequency (s = 1), but high restrictions on total retweet volume (n = 10,000, Figure 4.3). In this and future visualizations, we use the ForceAtlas2 (Jacomy et al., 2014) visualization algorithm as implemented in the application Gephi (Bastian, Heymann, and Jacomy, 2009). Intuitively, this means that two nodes are connected if at least 10,000 users retweeted both at least once during the period of this dataset, and nodes with more users retweeting both of them are more tightly linked together. Visualizing this projection reveals two tightly interconnected communities of influential users discussing the US presidential election, which we term and briefly describe as pro-Trump, and pro-Biden clusters. The pro-Trump cluster is anchored around Donald Trump's account, and includes an array of pro-Trump political activists, political organizations, politicians, media outlets, journalists, anonymous and self-identified online influencers, activists from the antifeminist "manosphere" online culture, and conspiracy-based QAnon communities. The pro-Biden cluster includes an array of politicians, journalists, media outlets, and online influencers, some of which self-identify as pro-Biden or liberal, and others of which, such as the television network CNN, identify as non-partisan but rejected false pro-Trump claims of election fraud.

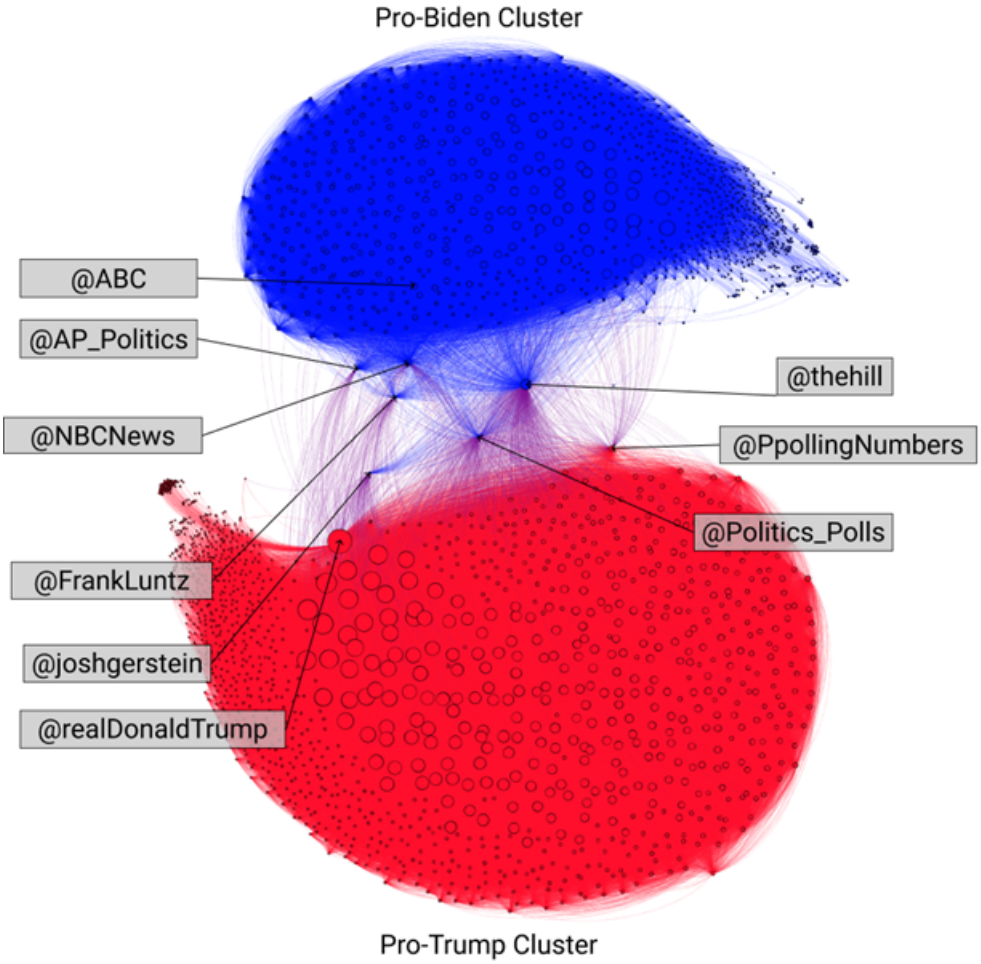

FIGURE 4.3: Case 1, election discourse for large audiences (s=1, n = 10,000). A co-engagement visualization of retweet relationships in a collection of tweets related to the US presidential election. Each node represents a Twitter account, and two nodes are linked if at least 10,000 users retweeted them both nodes at least once. Edges are undirected and weighted according to how many users retweeted both nodes. Nodes are sized according to their weighted degree, i.e. the sum of the weights of their incoming edges. Highlighted nodes represent nodes with connections to both pro-Trump and pro-Biden clusters.

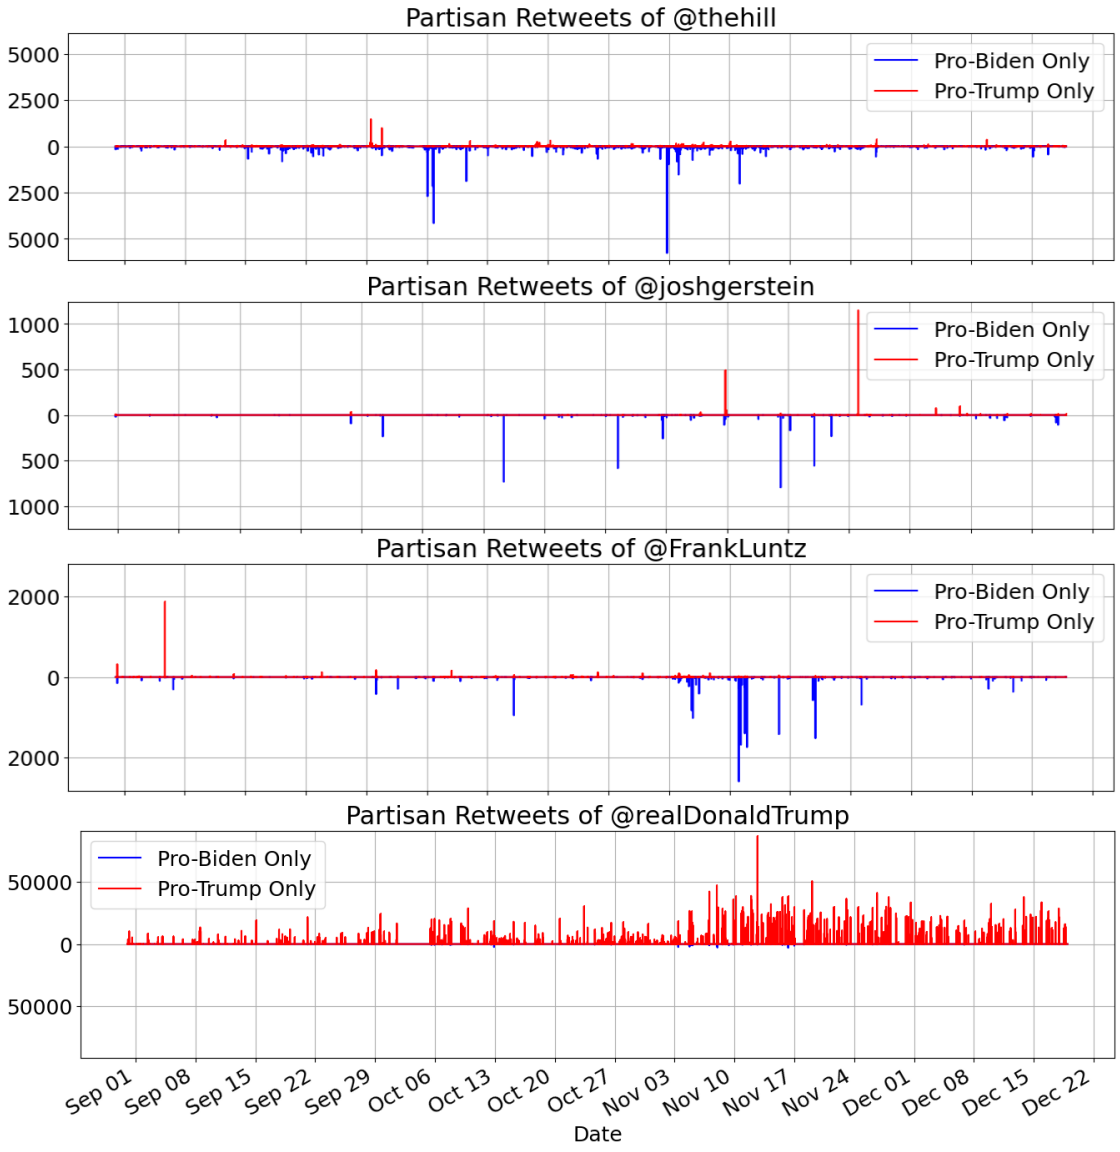

FIGURE 4.4: Example accounts with bipartisan engagement in Case 1. Plots of partisan retweets of four different bridging accounts labeled in Figure 3. Tweet dates are marked by the first time they were retweeted in our dataset. "Pro-Trump Only" partisan retweets indicate users who only tweeted accounts labeled in the pro-Trump cluster in Figure 3. "Pro-Biden Only" signifies the same totals for the pro-Biden cluster. Retweets of bridging nodes themselves are excluded in both retweet totals.

Overall, 2,499 nodes are generated in this graph, with 1,385 nodes in the pro-Trump cluster (55%) and 1,114 nodes in the pro-Biden cluster (45%). However, this set of retweets may be biased towards pro-Trump accounts due to our focus on terms related to election misinformation disproportionately spread by pro-Trump accounts.

We use the "pro-" descriptor to describe clusters as a whole, but some individual members of these clusters may identify themselves otherwise. We also use the terms pro-Trump and pro-Biden, rather than Republican and Democrat, to illustrate the extent to which traditional US party alignments are contradicted in the membership of these clusters. The Lincoln Project, a Republican advocacy group that supports conservative causes, is one of the most prominent accounts in the pro-Biden cluster, while many accounts in the pro-Trump cluster have ambivalent stances towards the Republican party outside of Trump. Their membership in a community is contingent not on their core beliefs or associations, but on their behavior in the dataset we observe, namely their posting behavior in the months before and after the presidential election. In datasets with the same users but different topics of discussion, the membership of individual users may change.

Audiences rarely retweet across clusters in large numbers, and when they do, they tend to retweet a select few cluster-spanning nodes. Almost all (98%) cross-cluster connections route through only eight nodes, labelled in Figure 4.3. These bridging nodes serve different functions in this discourse environment, and connect people in different ways (Figure 4.4). The most common form of bridging node was one where the account generally created two types of tweets, one that appealed to pro-Trump audiences, and one that appealed to pro-Biden audiences. The most transparent examples of these accounts were those that tracked polling results (@PpollingResults, @Politics_Polls, @AP_Politics, @NBCNews, @APPolitics), where those results and polls that favored Biden were retweeted by pro-Biden accounts, and those that favored Trump were retweeted by pro-Trump accounts. However, this form of apparent bridging also occurred when media accounts reported in neutral tones on events that fed preexisting pro-Trump and pro-Biden narratives respectively. The journalist account @joshgerstein was separately retweeted by both clusters for neutrally reporting

on Trump's attempts to contest the election results, with apparent pro-Trump legal judgments being more retweeted by pro-Trump users and their subsequent legal refutations more retweeted by pro-Biden users. The alternating quality of these nodes complicates the notion that they bridge communities, as their tweets are most often disproportionately shown to only one community at a time.

A special form of this alternating bridging occurred with the account @FrankLuntz, which posted updates on polls and predictions about the outcome of the presidential election. Before election day, this account was sometimes critical of Biden and released some predictions favorable to Trump, which led it to garner a slightly right-leaning cumulative audience. After the election, this account was resolute in affirming Biden's victory in the face of false pro-Trump claims of election fraud, then earning a growing pro-Biden audience. This account displays the sensitivity of such analyses to dataset selection, as a pre-election discourse analysis would likely place the account firmly in the pro-Trump cluster, whereas a post-election discourse in the pro-Biden, and when combined, firmly between. It was rare when accounts created posts that consistently appealed to audiences in both clusters. Some individual posts had equal appeal across clusters, such as when polling accounts released vote tallies tied at nearly 50% in critical states. Other posts that had a similar appeal were simply neutral statements of fact about recent news relating to the presidential election, which were made particularly often by the account for the online news organization The Hill (@thehill). While pro-Biden leaning, the Hill's account was one of the only accounts to consistently find engagement from both pro-Biden and pro-Trump accounts across this election time period.

The last significant point of crossover between the pro-Trump and pro-Biden clusters is the account for former president Trump itself (@realDonaldTrump). This circumstance reveals that though we imply for much of this analysis that retweets constitute endorsements, they do not always behave as such. Many pro-Biden accounts may be retweeting Trump's account simply

Top Election Accounts by Cross-Cluster Connections (Case 2)					
Biden-Trump		Biden-socialist		Trump-socialist	
Account Name	Share of Cross-Cluster Connections (n = 2,515)	Account Name	Share of Cross-Cluster Connections (n = 299)	Account Name	Share of Cross-Cluster Connections (n = 34)
thehill	24%	Proudsocialist	16%	ggreenwald	79%
PpollingNumbers	18%	BernieSanders	13%	jimmy_dore	9%
Politics_Polls	11%	davidsirota	12%	TulsiGabbard	6%
realDonaldTrump	7%	AOC	11%	aaronjmate	6%
AP_Politics	6%	briebriejoy	10%		
threadreaderapp	5%	KyleKulinski	7%		
Garrett_Archer	5%	IlhanMN	6%		
spectatorindex	4%	ryangrim	5%		
FrankLuntz	3%	peterdaou	5%		
DecisionDeskHQ	3%	RBReich	3%		

TABLE 4.1: Bridging connections between pro-Biden, pro-Trump and pro-socialist clusters. The top 10 Twitter accounts in terms of total cross-cluster edges for each two-cluster pairing in Case 2. These accounts share audiences of at least 100 users retweeting at least 5 times each across election discourse clusters.

because his tweets are often consequential in and of themselves, even when they are clearly opposed to Biden's election to the presidency. In several instances, pro-Biden users likely retweeted Trump's tweets sarcastically, such as when pro-Biden users disproportionately retweeted an old tweet from 2012 reading: "Scary thought–@JoeBiden is a heartbeat away from the Presidency."

4.6.2 Case 2: Third Party and Satellite Audiences at Lower Audience Sizes

In Case 1, a high n parameter demonstrated the nodes and graph structure of accounts with a relatively high volume of retweets over this dataset. While revealing the relatively few shared points of reference between pro-Biden and pro-Trump audiences, groups of users with smaller retweet bases are not visible in this graph. To visualize these communities, we generate the same data at a dataset with lower n, instead filtering by s to keep the size of the node and edge set tractable. Specifically, we generate a network in which links are defined when two nodes are shared by 100 users four times (n = 100, s = 5, Figure 4.5). After removing non-US clusters, this reveals a third group of nodes that we term the pro-socialist cluster. The pro-socialist cluster, much smaller than

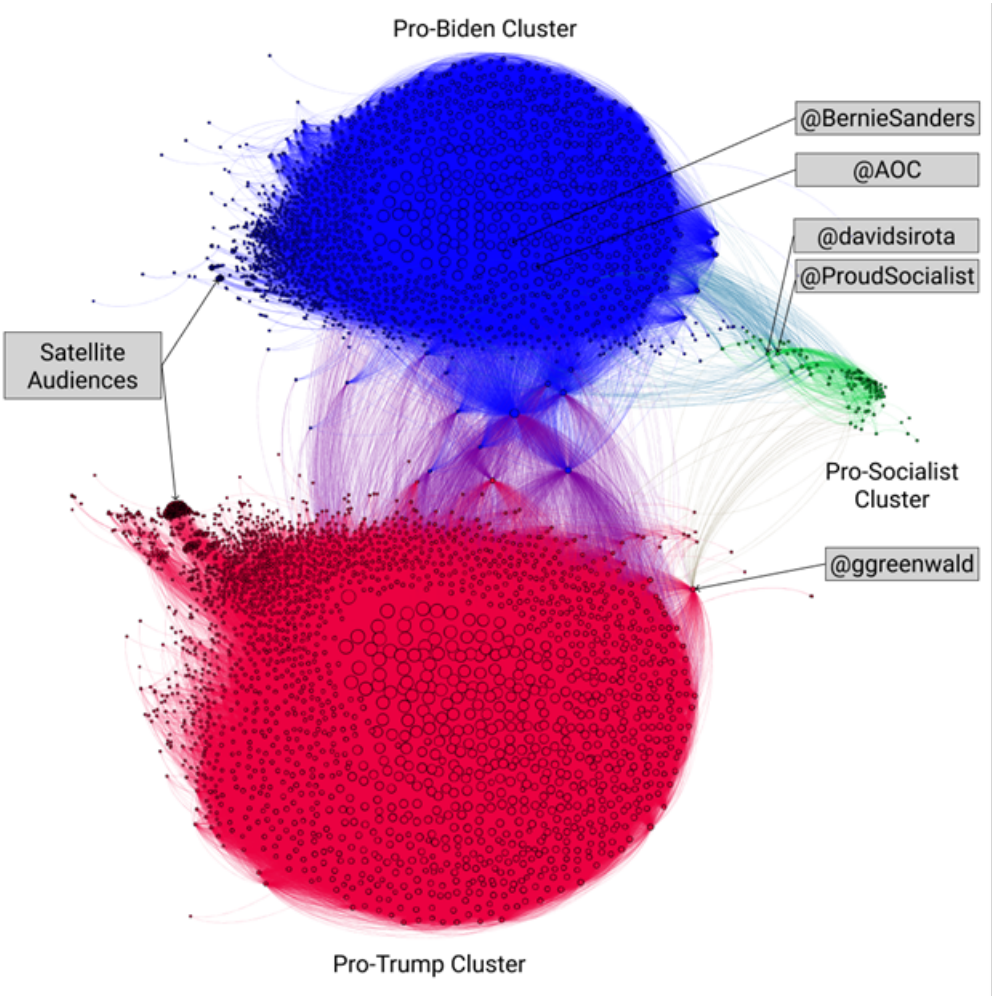

FIGURE 4.5: Case 2, Election discourse for medium audience sizes ($s = 5$, $n = 100$). A coengagement visualization of retweet relationships in a collection of tweets related to the US presidential election. Each node represents a Twitter account, and two nodes are linked if at least 100 users retweeted them both nodes at least five times. Edges are undirected and weighted according to how many users retweeted both nodes. Nodes are sized according to their weighted degree, i.e. the sum of the weights of their incoming edges. Nodes bridging between the pro-socialist and other clusters are labelled, as well as what we term satellite audiences: clusters of nodes with low-degree representing audiences with only tangential connection to mainstream US election discourse.

either the pro-Biden or pro-Trump clusters, is centered around multiple political activists, journalists, and influencers associated with US democratic socialist candidates and causes. We note that, like the pro-Biden and pro-Trump cluster, that this cluster does not contain all pro-socialist accounts, and some of its members may not identify as such. Indeed, popular democratic socialist presidential candidate Bernie Sanders is located in the pro-Biden cluster rather than pro-socialist cluster, likely due to his widespread popularity and continued public support for Biden during this election.

At a lower volume of retweeting, more bridges appear between the pro-Biden and pro-Trump clusters, and new bridges are generated between the pro-socialist and other clusters. Notably, the relative importance of some bridges changes under the new requirement for repeated engagement ($s = 5$). For example, the account for reporter @joshgerstein has no cross-cutting connections in this graph, despite being responsible for 7% of all connections in the first graph. This discrepancy is likely caused by how much of this account's pro-Trump retweet engagement comes from only two tweets, both, describing in a neutral tone, updates on pro-Trump attempts to legally invalidate the results of the election. This change demonstrates the effect of the s parameter, which can be tuned upward to select nodes for sustained engagement over a dataset, rather than widespread but momentary engagement in critical posts. The opposite effect can also be seen in those new nodes that appear as significant bridges. The spam media account @spectatorindex becomes a bridging node in Case 2, due to its sustained but low level of engagement, as do circumstantially important accounts like the governor of Arizona's (@dougducey), whose state was the center of many false election fraud claims.

Bridges between the pro-socialist and other clusters illustrate the different modes of inter-cluster commonality that can exist between different user clusters in this graph (Table 1). Between the pro-Biden and pro-socialist clusters, there are many nodes that draw consistent engagement, including popular democratic socialist politicians Bernie Sanders, Alexandria Ocasio-Cortez, and Ilhan Omar, and an array of writers, podcast hosts and other online influencers associated with

the socialist movement. Between the pro-Trump and pro-socialist clusters, however, there is little engagement, and all of it is mediated through three nodes. Most prominently among these nodes is Glenn Greenwald, a journalist who combined left-leaning views on topics such as surveillance with frequent engagement with right-wing media and criticism of mainstream media.

This second case study introduces a visualization feature which we had previously aimed to eliminate: tightly-packed clusters of low-degree nodes, in this case one-degree nodes mostly connected only to @realDonaldTrump. We note that while in ordinary graphs such nodes are usually uninteresting, in coengagement visualizations the relative isolation of these nodes reveals an important function in the election discourse environment. Particularly, many of these one-degree nodes are from communities plausibly isolated from mainstream US political discourse, but still displaying, for example, a contextual support for Trump or Biden. We term edges stemming from these nodes as *satellite audiences*, with inspiration from (Squires, 2002).

These nodes include accounts from other countries and/or in other languages, such as high-follower right-wing accounts writing for Japanese or Brazilian audiences. Such accounts are unlikely to interact with the majority of English-language right-wing accounts prominent in this graph, but may retweet Trump as a signal of nominal allegiance to his movement. Other low-degree nodes may originate in popular English-language, US-based communities that focus on topics usually unrelated to electoral politics. For example, the low-degree account for pop musician Ariana Grande, one of the most followed accounts on Twitter, connects only to Biden and fellow musician Lady Gaga, signaling a possible separation between entertainment-focused audiences and mainstream election-focused audiences.

4.6.3 Case 3: Followback Clusters at High-Frequency Engagement

Implicitly in the previous cases, structure is mostly determined by the number of users choosing to retweet two different accounts. However, as s increases and n decreases, structure is increasingly determined by repeated interactions by relatively small groups of users, which makes the actions

of well-coordinated groups more salient. To illustrate this, we generate a graph where links are defined by 25 users retweeting two nodes at least 25 times each over the course of the dataset (n = 25, s = 25, Figure 4.6).

In this case, three new clusters emerge with ties to the existing pro-Trump and pro-Biden clusters. We term the new clusters in this graph *followback* communities, due to their unique method of gaining followers and using Twitter. Accounts in these communities attempt to gain followers by mass-following other accounts in expectation of reciprocal follows, and sometimes explicitly coordinate with other accounts to expose themselves to a wider audience of potential followers. Because Twitter limits the number of users an account can follow by that account's current follower number, these accounts can often be distinguished from others by their nearly 1:1 ratio between followers and following totals (Figure 4.7). In addition to this follower manipulation practice, these communities have other unique behaviors compared to the clusters previously identified. Their median retweet total is much higher than that of the previously-identified clusters, and retweets are a much higher percentage of their total tweeting behavior. Their frequent retweeting likely propels their visibility in this visualization. Qualitatively, their behavior is also different from other users on the platform. They engage in retweet "trains,' in which they make posts tagging members of their own community and then retweet these posts incessantly in an attempt to garner more followers for all participants (Gallagher, 2020). They are almost entirely pseudonymous, with screen names and profile information often detached from any offline presence. One of the followback clusters is much larger than the others and associated with the core pro-Trump cluster, one is smaller and an off-shoot of the larger pro-Trump followback cluster, and the smallest cluster is associated with the core pro-Biden cluster.

As with previous clusters, we can investigate the nodes which bridge one cluster to another. In this case, however, both followback clusters have no connections to non-followback clusters not

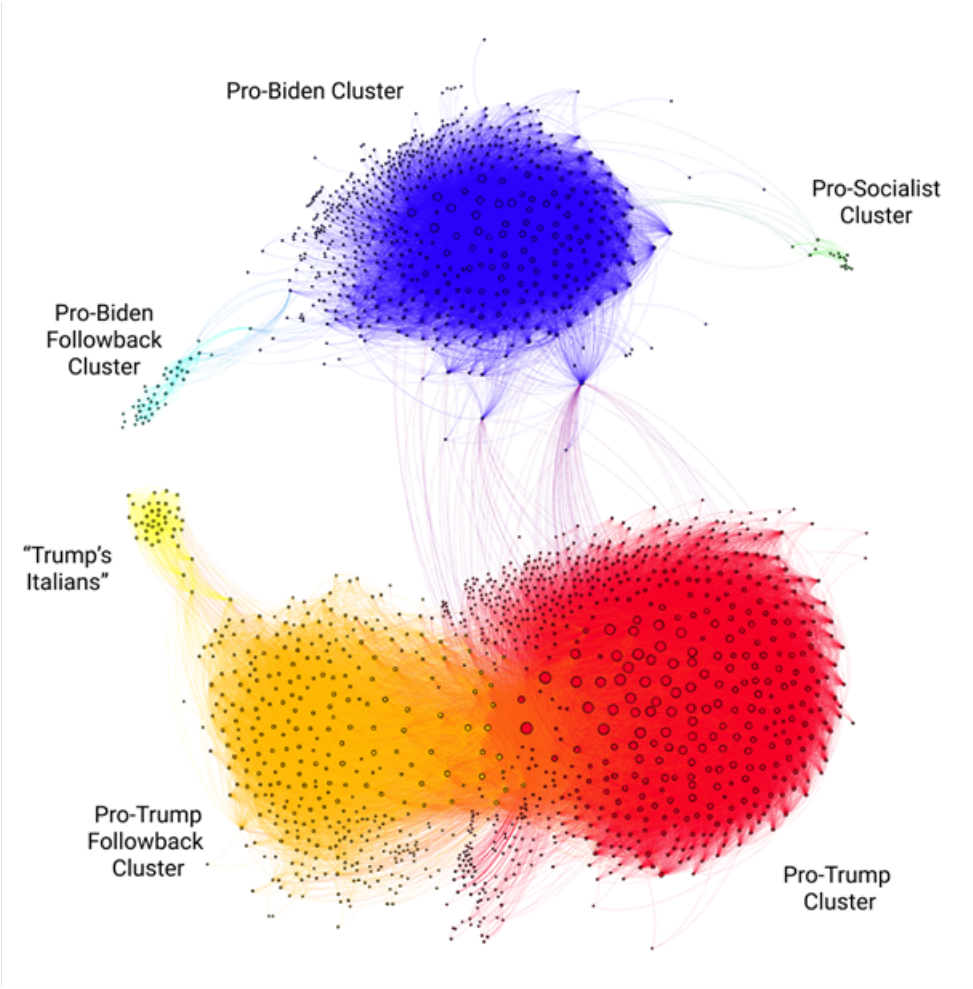

FIGURE 4.6: Case 3, Election discourse for small but active audiences ($s = 25$, $n = 25$). A coengagement visualization of retweet relationships in a collection of tweets related to the US presidential election. Each node represents a Twitter account, and two nodes are linked if at least 25 users retweeted them both nodes at least 25 times. Edges are undirected and weighted according to how many users retweeted both nodes. Nodes are sized according to their weighted degree, i.e. the sum of the weights of their incoming edges. Followback clusters, comprising nodes which retweet and follow other accounts to a relatively extreme degree, are labeled (Pro-Biden Followback Cluster, Pro-Trump Followback Cluster, "Trump's Italians").

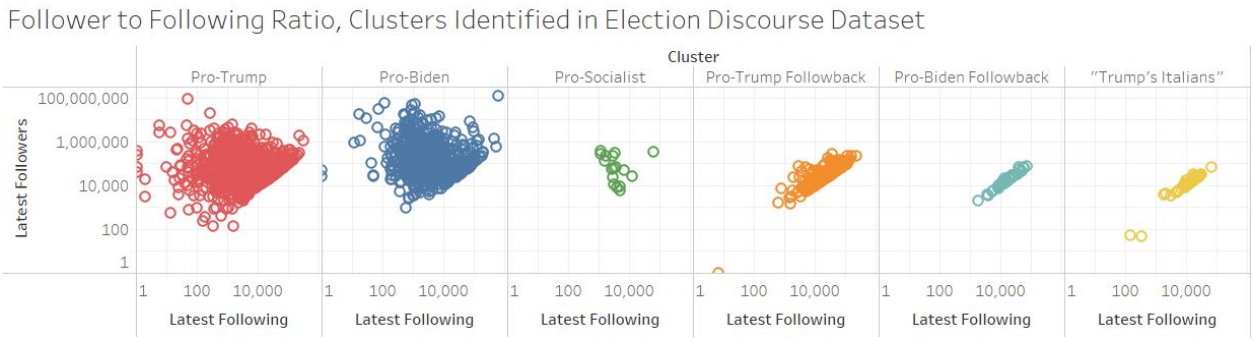

FIGURE 4.7: Follower to following ratios for Case 3 clusters. Scatter plots of the number of followers and the number of accounts followed for each Twitter account visualized in Figure 6 (n = 25, s = 25). Follower and following totals are counted at the time of the latest tweet recorded in the dataset. Plots are separated and colored according to cluster membership, and both axes are on logarithmic scale.

aligned with their preferred candidate. In all three cases, an important point of cross-cluster connection are the accounts for the two presidential candidates themselves (@JoeBiden and @realDonaldTrump), but unlike other groups, most nodes in the followback clusters are connected to these nodes. This feature of dense cross-cluster connectedness reflects a critical function of the engagement from followback clusters: to retweet the followback community, but also to retweet the influential nodes supporting their presidential candidate of choice.

Multiple followback clusters can exist supporting the same presidential candidate, and different clusters may share unique user characteristics. By inspecting the usernames and user-entered profile descriptions of the smaller pro-Trump followback cluster, we found that 29/41 accounts provided some indicator of Italian American identity, either explicitly identifying as such, identifying as part of "Trump's Italian Army," or including Italian flag emojis paired with US flag emojis. While the larger pro-Trump followback also contained some self-identified Italian users, they were by no means the majority as in the smaller cluster. This smaller cluster was linked to the main pro-Trump followback cluster only through four bridging nodes, two of which identified as part of Trump's Italian army, and two which self-described as duplicate accounts for the

same user. This cluster otherwise only formed internal connections and connections with @realDonaldTrump and @Llinwood, the account of L. Lin Wood, a lawyer who advanced many false conspiracy theories relating to the election and litigated on Trump's behalf in some post-election court cases.

All three followback clusters reflect the curious quality of coengagement networks where users can play both the role of edge and node. Recall that coengagement networks are graph projections, where each edge is comprised of engagements from many unseen nodes present in a ordinary, unprojected version of the graph. Almost 38% of the unprojected nodes that comprise the edges in followback clusters are themselves represented as nodes in the coengagement networks, compared to less than 1% of the accounts in the core pro-Trump, pro-Biden, and pro-socialist clusters. This is to say that followback clusters uniquely play the role of their own audience, relentlessly sharing their own members' content in addition to their chosen candidates. This feature sheds light on how the smaller pro-Trump followback cluster can separate itself in this graph, as its self-identified Italian American accounts specifically and at the threshold we set retweet only each other. This analysis does not allow us to speculate as to the level of explicit coordination or the motives behind these accounts for acting in this way. While some of these clusters' unusual behavior could indicate automated or semi-automated activities, previous reporting has indicated that at least some of these accounts are likely operated and coordinated by otherwise ordinary users that simply aim to support Trump's candidacy on Twitter (Gallagher, 2020).

4.6.4 Synthesis: Cluster Contingency and Differential Moderation

By examining the full *n/s* parameter space, we can develop a sense for the relative size and sharing characteristics of the clusters that we have identified so far. We calculate the results of repeated network clusterings using the Louvain algorithm on a range of possible n and s parameter combinations, assigning cluster labels via high-degree landmark nodes associated with each previously-identified cluster. In Figure 4.8, we illustrate the boundaries at which these clusters are

no longer salient. The pro-Trump, pro-Biden, and pro-socialist clusters are all contained within a maximum number of engaged users that steadily declines with restrictions on these users' number of retweets. By contrast, the followback clusters cannot be detected unless users' minimum retweets are elevated, but are never salient at a size of greater than 110 users. We note that the point at which a cluster fails to be detectable is not the point at which all nodes in these clusters are removed, but rather the point at which these nodes become subsumed into larger clusters. For example, the two highest degree pro-socialist nodes from Case 2 are relatively-marginal nodes in Case 1, while many of the most popular Trump followback accounts in Case 3 are found in the core pro-Trump cluster in Cases 1 and 2. These facts reiterate that the clusters we identify here are not found with respect to the actual ideologies or social ties between nodes, but rather through an understanding of how nodes are perceived by their engaged audience variously defined.

We conclude these case studies by noting how the clusters we identify here were subject to different levels of moderation in the wake of the US presidential election (Figure 4.9). On January 6, 2021, a large group of rioters, associated with a range of pro-Trump movements contesting the results of the 2020 presidential election, entered the US Capitol building while its representatives were in session. Following this event, Twitter suspended a large number of accounts said to have encouraged this violent protest, as well as accounts associated with the QAnon conspiracy movement (Romm and Dwoskin, 2021). To measure the effect of these and other more recent suspensions on our dataset, we identified all accounts that were suspended as of September 3, 2021. Since the election, most of the pro-Trump followback cluster in this case study (71%) has been suspended by Twitter, compared to 32% of accounts in the core pro-Trump clusters, 2% of the core pro-Biden cluster, 7% of the pro-Biden followback cluster, and 2% of the pro-socialist cluster. In other words, these suspensions have disproportionately affected the pro-Trump followback users we identified in Case 3, which is to say groups of pro-Trump users with frequent intracommunity retweet activity.

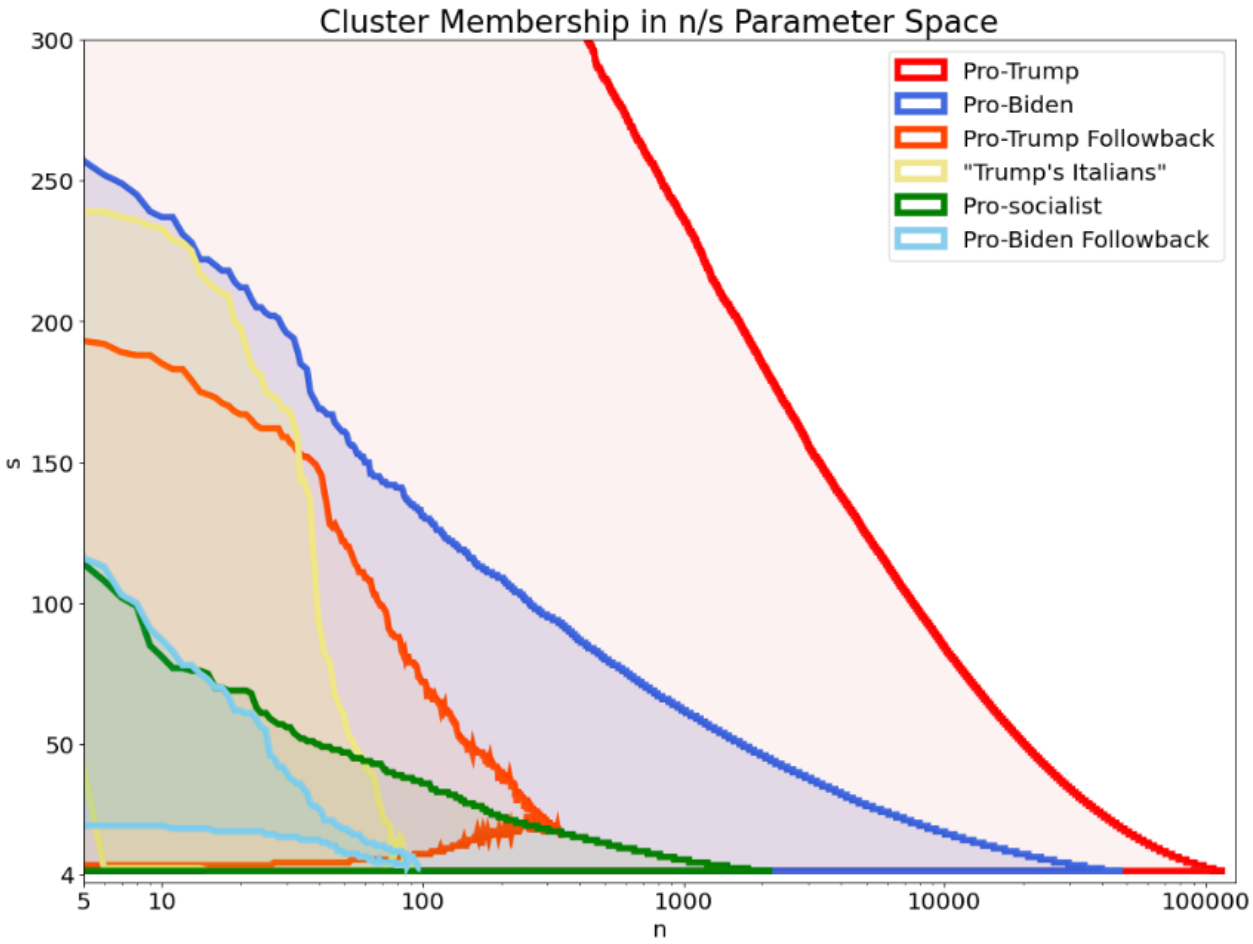

FIGURE 4.8: Existence map for clusters across filtering values n, s. For each cluster identified in these three case studies, we determine the maximum parameter values at which these clusters can be identified. Clusters are labeled at each (n, s) parameter value if they contain specific high-degree landmark nodes identified in the previous case studies. In the shaded regions, individual clusters are salient to the clustering algorithm. Outside of the shaded regions, these clusters cannot be identified either because their constituent nodes are not present or because they have been subsumed into other clusters.

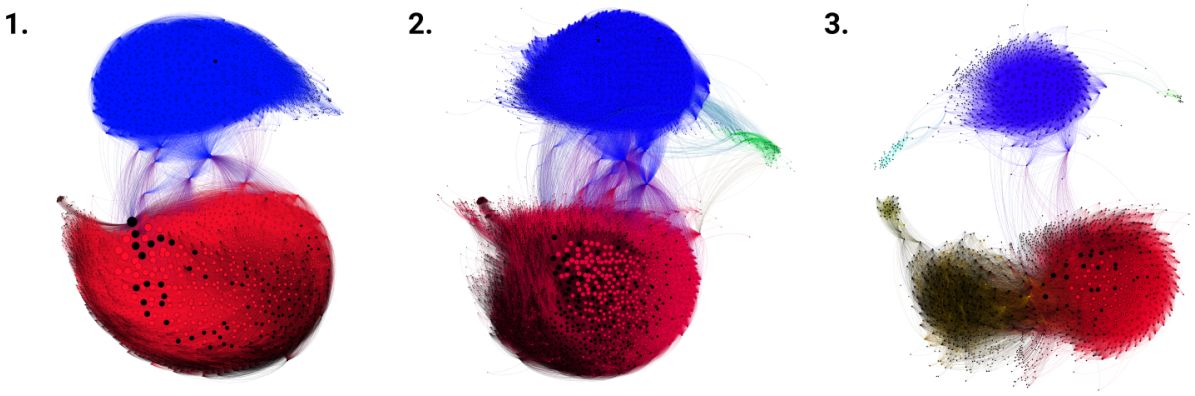

FIGURE 4.9: Visualization of suspensions across all case studies, labeled (1-3). The original networks from all three case studies are visualized in increasing order from left-to-right. Nodes that have been suspended by Twitter as of January 6, 2021 are colored in black, otherwise coloring remain the same as in the original case studies.

4.6.5 Comparison: Directed Engagement Graph and Flows

We conclude our case studies with a brief comparison to a more standard network form, which we refer to as the directed engagement graph. In directed graphs, each node is an account, each edge is a directed retweeting relationship, and edge weights signify the frequency of engagement. In this present dataset, the unfiltered directed graph consists of 22M nodes and 299M edges. This graph cannot be visualized effectively with our current software and computing power, although a version with heavily filtered edge weights can be observed in Figure 1.

We first assess whether coengagement graphs represent nodes typically perceived to be important in the directed graph. We compare the weighted (in)degrees of both graph, which in the case of the directed engagement graph is simply equal to the number of times an account has been retweeted. We find that of the top 1000 most retweeted accounts, 95%, 96% and 85% are represented across Cases 1, 2, and 3 respectively. Furthermore, we find that the nodes we choose to include in Cases 1, 2, and 3 account for 54%, 64%, and 51% of *all* retweets in this dataset, despite only representing fewer than 0.1% of the accounts in our data. There are some highly-retweeted nodes in the directed graph that are not visible in each coengagement network. However, these missing nodes are also

mostly unrelated to the 2020 US presidential election. For example, the most retweeted nodes missing from Cases 1, 2, and 3 respectively are a fan account for pop musician Justin Beiber, a fan account for the Korean pop music group BTS, and the account for Billboard, a music media outlet. These nodes are highly retweeted in these datasets because of competitions in which fans "vote" for their favorite musicians. They may be missing from our coengagement visualization simply because they are essentially apart from the apparent main topic of the dataset, and thus have fewer opportunities to "share" an audience with another node that is collected under these terms. Some relevant highly-retweeted accounts, such as the Twitter account for Donald Trump's daughter Ivanka Trump, are also missing from Case 3 likely because they did not publish enough election-related tweets to reach the 25-retweet threshold.

We second assess whether structures found in the directed graph are significantly different from the coengagement graph. To do this, we cluster the directed graph using the Infomap algorithm (Rosvall, Axelsson, and Bergstrom, 2009), a different clustering method which views directed, weighted edge interactions as flows of information between nodes. We find that the pro-Biden, pro-Trump, socialist, and pro-Trump followback clusters are all still found under this algorithm, with the pro-Biden followback cluster being subsumed into the pro-Biden cluster. No other large cluster is found which combines nodes found in these case studies into new mixtures. We do find, however, that clusters apparent in these case studies, such as the pro-Biden cluster, often are subdivided into smaller clusters in the directed network.

These smaller clusters reveal the tradeoffs between directed networks and coengagement networks, as they most likely stem from Infomap's tendency to privilege connections between high-indegree nodes, i.e. retweets between influential accounts. For example, a pro-Biden subcluster in the directed graph is centered around official accounts and reporters for the media outlet The New York Times. These popular accounts frequently retweet each other, probably to promote each others' work, which strengthens their association in graph forms and algorithms which understand high-degree nodes as routes through which users travel. By contrast, in coengagement graphs,

the influence of interactions between influential nodes on the resulting form and clusters is dampened, due to its focus on *large* but not necessarily *important* shared audiences. In social engagement data in which user engagements can be seen as travelling from node to node, such as use clicks through profile pages, directed retweet networks and clustering algorithms that consider them may be more appropriate. In data such as this Twitter retweet dataset, where engagements are enacted from relatively stable positions and circumstances, coengagement graphs may be more appropriate.

4.7 Discussion

We introduce coengagement networks and illustrate their value for a mixed methods analysis of a dataset of Twitter posts related to the 2020 US presidential election. We illustrate how the number of apparent clusters perceived in these networks is contingent on the minimum proposed size and activity level of their engaged audiences. When seen through the lens of large, momentarily interested audiences, there appears to be two dominant pro-Trump and pro-Biden communities. When including smaller engaged audiences, a pro-socialist cluster emerges, and when focused on highly-active but even smaller engaged audiences, unique followback clusters emerge with severely different user behavior. These networks make clear that Twitter's moderation in the wake of the attack on the US Capitol Building disproportionately affected these pro-Trump followback accounts, while also affecting a number of pro-Trump accounts with more ordinary retweeting behavior. Taken together, the insights from these networks depict an ecosystem of popular and activist discourse communities in the presidential election with few but crucial points of overlap, and the *de facto* removal of the majority of influential member is one of these communities in the wake of the US Capitol attacks.

The purpose of our case study is not necessarily to draw definitive conclusions or cause-and-effect relationships, but to use visualization to expose those features of this discourse ecosystem that deserve further study. We describe a mostly binary structure of pro-Trump and pro-Biden

engagement in English-language US election discourse, and provide a typification of the points of crossover between these two clusters of accounts. We identify a detectable and yet marginal tide of third-party US political discourse in the pro-socialist cluster whose growth and comparative influence on non-election discourses may yet have further importance. We identify the phenomenon of satellite audiences, where high-degree nodes, and particularly the account @realDonaldTrump, serve as singular points of reference for many communities with only marginal connection to English-language election discourse. And we characterize followback communities, which use unique posting strategies and engage in unusually partisan rhetoric to support specific candidates in our election dataset.

The social network visualization approach described here makes transparent many features which contribute to this understanding of Twitter discourse around this election, while reducing visual clutter and artifacts typical to equivalent large datasets. However, we stress that the sum total narrative could not be found from any of these visualizations taken alone, and indeed some visualizations have contradictory features which could, when viewed in isolation, generate misleading insights as to the nature of these communities. For example, we have shown that the pro-Trump followback cluster is from the perspective of the size of their engaged audience a somewhat marginal phenomenon, ceasing to be organizationally coherent at the level of more than a hundred users. Portraying them with equal prominence to more typical political communities may mislead viewers as to their relative impact on the overall election conversation compared to, for example, the core pro-Trump cluster. Yet small groups of coordinated users can nonetheless have an outsize impact in online communities (Center for Countering Digital Hate, 2021), and our analysis of Twitter suspensions since the presidential election show that this followback community was specifically targeted for moderation at a much greater rate than other communities. Depending on the visualization goals of the researcher — proportional impact, or behavioral diversity — forefronting this community may or may not be informative to their chosen audiences.

Portraying multiple visualizations that provide both interpretations, and analyzing the discrepancies between them, provides a level of explanatory power that no single visualization is likely to otherwise achieve.

Overall, what we describe in this paper is an interpretative, mixed-methods workflow, in which visual artifacts derived from quantitative network transformations are combined with a deep qualitative understanding of the US Twitter election context, to the benefit of both. Given the contingency of network visualizations on their initial parameters, deep understanding of the context of a given dataset is necessary to interpret and triangulate their different incarnations. However, given the size of contemporary social interaction datasets, quantitative methods such as the coengagement network are necessary to reduce the complexity of online conversations to artifacts which are tractable to qualitative researchers. We advocate for the continued use of coengagement networks in mixed-methods research, as a tool that both stimulates and benefits from deep contextual understanding of increasingly unwieldy social data.

4.7.1 Ethical Considerations, Limitations, and Software Sharing

As network visualizations continue to be central in social media analysis, we believe it is necessary to briefly examine the ethical considerations on whether network visualizations such as these *should* be used in every circumstance. We have chosen here to visualize users participating in a high-prominence topic, consisting of mostly public-facing accounts such as politicians and media outlets. The same methods applied to communities with a higher expectation of privacy, or who face higher risks from exposure, may be unethical surveillance if researchers have not derived consent from members of these communities. There are also ethical implications to naming accounts visualized as nodes in networks. Some users, due to gender, race, or other factors, are at higher risk of harassment if identified as influential in a given community, while other users who explicitly seek attention in online communities may use their identification in networks as a propaganda tool in hateful campaigns. In our reporting of this work, we have declined to name

some accounts for both reasons. We also stress here how the data collection procedure, and subsequent description of that procedure, affects which communities appear to be participating in a phenomenon. Several politically-active Twitter communities in the US that have been previously described in research, such as Black Twitter (Clark, 2019) and non-English language communities (Soto-Vásquez et al., 2020; Fang, 2021), are not explicitly visible in our analysis, likely due to their different posting volumes and the choice of terms and topics on which we chose to center our data collection. Researchers working with such visualizations whose research bears on policy and public perception must explain such limitations in the communication of their work.

We have described the basic form of an approach for visualizing engagements in social network data, and there are many ways in which this method can be modified to more saliently capture engagement dynamics. For example, the current formulation of this method places emphasis on users who frequently share content, which is not necessarily undesirable given the external impact of this behavior. However, different formulations of the network projection scheme, such as those that weight users' engagements relative to their average level of engagement, may be a better reflection of real discourse communities that exist at lower sharing volumes. We also observe that many of these coengagement networks create densely-connected subgraphs in which most nodes are connected to most other nodes, making internode relationships difficult to visually identify. Accordingly, these networks may be complementary with other techniques to improve visualizations of social networks, such as the edge sparsification procedures for densely-connected networks proposed by Nocaj et al. (2015).

In the hopes that others may replicate our methods of both visualization and analysis on new datasets, we make the code available for generating these graphs either from structured data received from the Twitter API, or in general JSON and CSV-based formats.[1] This code uses the visualization capabilities of the open-source network visualization library Gephi, and its implementation of the ForceAtlas2 algorithm for network visualization (Bastian, Heymann, and Jacomy,

[1] https://github.com/uwcip-research/Coengagement-Networks

2009; Jacomy et al., 2014). We have packaged this code in publicly-available Docker containers, a relatively portable and stable code format which can be run on many machines with relatively few installation requirements. Additionally, we have made available node and link data for all visualizations displayed in this paper, as well as a list of Twitter ID numbers for tweets and users corresponding to data used to generate these graphs. We hope by making the code for generating these graphs open-source, other researchers both qualitative and quantitative will both explore the potential and limitations of this method, as well as contribute modifications to this scheme as appropriate.

4.8 Acknowledgments and Funding

We would like to acknowledge Clément Levallois for his advice on the open-source implementation of this project, and Paul Lockaby for ample support in data collection. We received funding from the University of Washington's The Center for an Informed Public, The John S and James L Knight Foundation (G-2019-58788), Craig Newmark Philanthropies, and the Omidyar Network.

Chapter 5

Study 2: On Measuring Change in Networked Publics

Contributing Author: Anna Beers[a]

[a]University of Washington, Human Centered Design and Engineering

5.1 Preamble

This chapter is currently under revision at a peer-reviewed journal, and the text and figures are only slightly modified from that original submission. This preamble serves as additional contextualization for this study within the context of this dissertation, and personal reflection on my part for where this work has since travelled. I use single-spacing in this preamble to offset it from the main text.

Like Study 1, this study was inspired by work done for the University of Washington's Election Integrity Partnership, this time assembled for the 2022 election. It was driven by my intuition that right-wing influencer Twitter activity during the 2022 election seemed unusually subdued relative to 2020. But how would I even know, given the multitude of factors at play, the differing meanings of the word "change," and my limited perspective? This study was then an attempt to tackle this question head-on. Within the context of the dissertation, this was a way of measuring and representing influencer community structure through the lens of change: which parts change, and how could we measure that change.

As the study describes, understanding if and how much influencer communities have changed can inform public and policy discourse in productive ways. For me, this understanding was

somewhat personal: how could I evaluate my and my colleagues work in the 2020 Election Integrity Partnership? While we did not intend to intervene directly in election publics, our work highlighting the spread of misinformation among right-leaning publics, combined with members of those publics later taking offline action at the U.S. Capitol building, likely informed subsequent policy which trended towards their increased moderation.

This study, of course, does not answer whether that moderation resulted in some of the diminished right-leaning activity we see here. Personally, I struggle to visualize any study which could test for such a proposition. Moderation, one assumes, has complex effects on influencer activity. Some influencers are removed from the platform, but other influencers likely take their place. However, the influencers which take their place likely adapt to the signals given off by the prior influencers' dismissal, and moderate their own language accordingly. New and remaining influencers may simply not engage in topics which could one day subject them to removal, seeming to decrease their presence *only on that topic.* On top of all that, non-platform factors, such as in this case the possibility that right-leaning influencers determine online rhetoric was not terribly effective in winning real-world elections, could also contribute to changes in the appearance of right-leaning activity. While this study offers tentative conclusions about change in influencer communities, it mostly chooses to embrace the ambiguity.

Something which ultimately fascinated me the most about this study was the fact that while our ability to estimate how communities "change" is weak, many of us attempt to make such estimates anyway. Though I do not explore it here, I was particularly compelled by how users and influencers themselves might be understanding changes in their own communities. While I, as the comparatively well-resourced academic, have access to more data than they do, they have access to a multi-dimensional lived experience in this community which likely provides them with an even-clearer picture. As scholars of influencers often remark, they watch metrics the same way that I do, and make decisions accordingly (Hund, 2023, among many others). It would be interesting to simply interview the influencers involved in this study for their recollections of how they believed their networks have changed. However, there would be no reason to trust their reflections directly, given that an opportunity to opine in a published academic article would itself be folded into their self-branding efforts; it can be beneficial to be seen as part of a growing community, or conversely, part of a persecuted one.

The use of "election publics" and literature review on communities, publics, and networks represents a step in my evolving understanding of influencer aggregates. My dissertation proposal originally used the phrase "influencer publics" as opposed to "influencer communities," and I think there are compelling cases for both terms (as well as others). Discussing publics of influencers connects them with a greater literature on publics, which has self-consciously analyzed the sorts of political communities I profile in these case studies. However, influencer communities suggest a greater stress on the interpersonal nature of influencer interactions, the speciation of different influencer subtypes, and the ways in which the audience experience of influencers is often watching a *community* — what I sometimes describe to my colleagues as "the social network projected into the sky." I look forward to seeing how our field finds ways to talk about influencer aggregates in the future.

It should be briefly noted that this study was a test case of a sort of event-based longitudinal analysis now made impossible in the post-API era. One could imagine that long-running groups like the Election Integrity Partnership, collecting similar data on similar events at regular intervals, could start to produce valuable knowledge on these events from a longitudinal perspective at long enough time spans. This was always a bit of a pipe dream; the changing nature of platforms and tenuousness of our access to data always made it unlikely that academics would have comparable data across timepoints. However, I do think this tentative longitudinal analysis provides an interesting case study for more well-resourced organizations who *do* have the resources and ability in this manner.

5.2 Abstract

Networked publics are the imagined communities that form on technologically-mediated spaces, often to deliberate on current events and enduring issues. Measuring change in the size of networked publics is practically difficult because of the different ways that change in such communities can be operationalized. Yet measuring how networked publics are changing is particularly pressing, as researchers, platforms, and civil society have proposed interventions to curb their behaviors in the wake of digitally-mediated disinformation and violence. Here, I propose a trifactor model for measuring change in publics, and demonstrate this model on a case study of United States networked election publics on Twitter from 2020 and 2022. These three factors center on the creators of media, their audiences, and the digital platforms which mediate them, and analyze both the total membership and communicative activity of these publics. Analyzing over 600M posts from 2020 to 2022, I find that right-leaning creators and audiences decreased in their activity relative to left-leaning publics in 2022, and faced disproportionate platform moderation due to suspensions, removal of spam networks, and shadowbanning. However, these negative signs are offset by a relative growth in the size of right-leaning publics due to higher user entry, and a higher volume on non-election topics. These changes shed light on a widely-publicized effort to moderate disinformation and incitements to violence in right-leaning election publics on Twitter in the wake of the 2020 U.S. presidential election.

5.3 Introduction

Networked publics are the imagined communities that form in technologically-mediated spaces, often to deliberate on current events and persistent issues (Boyd, 2007). Measuring change in networked publics is difficult, and yet much popular discourse about the internet is seemingly predicated on our ability to understand these changes. Take, for example, the widely-publicized disinformation and incitements to violence that increasingly occur in networked publics during United States (US) national elections (Center for an Informed Public et al., 2021; Donovan, Fagan, and Lee, 2022; Chen, Lukito, and Koo, 2023). Researchers, journalists, and politicians have sought interventions which would moderate the spread of misinformation on these networked election publics (Ovide, 2021), but there are few obvious, public-facing indicators for designers of interventions to understand how these interventions affect the longitudinal development of these publics as a whole. Efforts to create such indicators are complicated because data about the size and activity of networked publics is restricted and usually partial, and regardless qualities like *size*, *activity*, and even the contours of *publics* themselves are always defined multiply and contingently.

In this paper, I present a multi-factor model for understanding change in networked publics, and apply this model to a case study of change in the wake of misinformation and incitements to violence in partisan Twitter publics from the 2020 United States presidential elections to the 2022 United States congressional elections. I come to three main conclusions. First, concerning the numerical membership of these publics, right-leaning participation in election publics increased in membership relative to left from 2020 to 2022, even in the face of disproportionate platform suspension rates caused by their members participation in election fraud narratives post-2020 (Romm and Dwoskin, 2021). Second, despite this growth, right-leaning creators on average made relatively fewer election-related posts than left-leaning in 2022, and right-leaning audiences were similarly less likely to share those posts. The activity of right-leaning spam networks, previously implicated in sharing low-quality news and conspiracy theories (Torres-Lugo, Yang, and Menczer, 2022) was particularly decreased in 2022 when compared to similar left-leaning networks. Third,

concerning the ability of users to reach their audiences, I show limited data suggesting that right-leaning elite users, many of whom were were found to have spread election misinformation during the 2020 election (Kennedy et al., 2022), were more likely to be shadowbanned than left-leaning users after the 2022 election, and that their follower counts grew more slowly than left-leaning accounts since 2020.

Overall this model suggests an analysis of change in the communicative output of networked publics via a three-way link of creators, audiences, and platforms. I discuss how changes in these groups combine to create external perceptions of publics' communicative output, and also how a change in one group is inextricably related to changes in another. In this particular case, I discuss how to interpret these different, conflicting indicators of change in light of highly-publicized suspensions of pro-Trump, conspiracy-minded users that occurred in the wake of the 2020 US election (Romm and Dwoskin, 2021; Booker, 2021). I stress the urgency of developing ways to measure changes in networked publics in an era in which social media sites face public pressure to moderate them. I also reflect on the difficulty of assigning cause and effect in internet publics when access to large-scale data is increasingly limited by private companies. Specifically, relying on imperfect metrics like the retweet to assess activity in publics has clear drawbacks, but such measures are also the most clearly visible and intuitive signs of a networked public's growth to civil society.

5.4 Background

The primary task in studies of large social media datasets is to determine a mapping from a dataset to a given conceptual object, usually representing an aggregation of people. Most studies correspond datasets to either the *online community*, *online social network* (*OSN*), or the *networked public*. The boundaries between communities, networks, and publics are not always clear (or even well-defined), and accordingly I briefly review prior literature on measuring change in each of these forms.

5.4.1 Change in Online Communities

Online communities are typically defined as emergent social aggregations of people who, through repeated interaction on the internet, develop strong relationships and a self-conscious "sense of community" (Rheingold, 1993). The study of change in online communities has typically been seen as a question of balancing platform costs and benefits to new and existing users (Garcia, Mavrodiev, and Schweitzer, 2013). Change is often benchmarked against a certain measure of community "success," marshalled against failure or death. Community success can be defined in different ways: the number of members, the frequency of interaction, the denseness of connection between actors, the quality of those connections, or the quality of output from that community (Iriberri and Leroy, 2009). Some apparent measures of success can be inversely correlated with others, such as the supposed decline in community standards and increased antisocial "troll" interactions that comes with influxes of newcomers (Kiene, Monroy-Hernández, and Hill, 2016; Cheng et al., 2017). Different users also contribute in different ways to the success of online communities, with particular attention being given to the distinction between participants, who generate interactions, and "lurkers," who passively observe those interactions (Malinen, 2015).

Generally, online communities are observed to follow a "lifecycle," where early periods of creative inception give way to growth, stability, and possibly demise – often through outmigration to other platforms (Iriberri and Leroy, 2009; Newell et al., 2016). Designers of social media platforms have taken particular interest in designing features which afford community growth and stall decline (Kraut and Resnick, 2012).

5.4.2 Change in Online Social Networks

Social media datasets are also studied more strictly as social networks, where internet users are nodes and their relationships or interactions are links. Online social networks (OSNs) may contain many online communities, but more closely resemble a structured medium upon which online communities are formed. The node and link schematic reduces and simplifies the ways an OSN

can change. Nodes and links can be added or removed from a graph, and edges can either be added, removed, or rewired from one node to another (Valente, 2012; Ghoshal, Chi, and Barabási, 2013). Most OSN researchers attempt to characterize internally-driven (endogenous) network *evolution*, whereby intrinsic properties of networks (e.g. reciprocity), or intrinsic properties of their actors (e.g. homophily), drive the adding and rewiring of links within networks over time (Stokman and Doreian, 1997; Carley, 1999; Kossinets and Watts, 2009; Block, 2015). Fewer network theoretic studies have addressed the effects of exogenous shocks and *initiated change* in networks, like content moderation and crisis events, that typically affect contemporary social networks (McCulloh and Carley, 2011).

5.4.3 Change in Networked Publics

Distinct from networks and communities is the public, for which there exist multiple and somewhat conflicting definitions. Most definitions have their root in the idea of the Habermasian public sphere, generally defined as a mediated space where a nation's citizens come together to deliberate upon "public" issues in a structured and rational manner with the goal of informing governing priorities (Habermas, 1991). Over time, this concept has been relaxed to include the existence of multiple and different types of public spheres, spanning different audiences and serving different, sometimes non-deliberative roles. (Fraser, 1990; Squires, 2002; Papacharissi, 2015). Supplementing the study of the public sphere is the concept of issue publics, originating separately with political scientist (Converse, 1964). An issue public is simply the aggregation of people interested in a certain issue, such as immigration or the access to abortion. Habermas later amended his theory of the public sphere to acknowledge the existence of issue publics, essentially combining these academic lineages into one meta-concept (Habermas, 2006). With internet media came the rise of different sub-terms to describe internet publics and public spheres, including the blogosphere (Tremayne et al., 2006), the Twittersphere (Bruns et al., 2017), Twitter publics (McKelvey, DiGrazia, and Rojas, 2014), and most relevantly here, networked publics (Boyd, 2007; Varnelis, 2008).

Networked publics share qualities of both online communities and OSNs. As defined by boyd, a networked public represents the technologically-mediated space upon which online publics form, similar to the broad definition of an OSN. However, per the same definition, a networked public also signifies the imagined communities which form upon that online strata, echoing studies of online communities. Studies of networked publics often focus on how the affordances of social media platforms shape possibility for online deliberation and creativity (boyd, 2011). The typical measurement of change in networked publics focuses on dynamics in the event-scale: how actors in publics collaborate or compete with each other to accomplish goals within a time-bounded debate. Examples include the study of mobilization for political action in networked publics (Benkler et al., 2015; Jackson and Foucault Welles, 2016), and the study how different (counter)publics compete to set political agendas (Freelon, McIlwain, and Clark, 2018; Yang et al., 2021). Longitudinal (multi-event) studies of networked publics are more rare, possibly because publics are often seen to be irreducibly tied to the single event or issue which called them into being.

5.4.4 Networked Election Publics

In this study, I consider the publics which develop in anticipation of government elections on social media to be a recurrent, periodic form of networked publics. Their recurrence at regular intervals, the persistence of political parties as proxies for online communities, and their organization around a single time-delineated event let researchers understand how ordinarily unique networked publics change in the face of exogenous shocks and longitudinal endogenous changes. Due to the rigorously structured affordances of Twitter and other social media platforms, these networked publics are easily comprehended as examples of OSNs: collections of nodes (users, posts) connected via links (shares, follows, likes). Elections as networked publics also feature the prominent involvement of sustained online communities, with organized and self-consciously affiliated groups of users associated with political parties marshalling their efforts to advocate for their candidates. Accordingly, I propose a model for modeling change in these publics that borrows from all three of these lineages, while synthesizing them into an object of study called the

networked election public.

Previous longitudinal studies of internet use during elections have focused on the total number of, frequency of, and distribution of online interactions within networked election publics, usually according to party affiliation (Larsson and Moe, 2015; Lilleker, Koc-Michalska, and Jackson, 2015; Bruns and Moon, 2018). Others have sought to expose trends in how online participation during elections have translated to offline activity, such as voting or self-rated political efficacy (Zhang, 2016; Jensen and Schwartz, 2021). Still others have assessed the changes in the content of election-related messaging, with respect to tone, topic, and external linking to news sources (Bruns and Moon, 2018; Stromer-Galley et al., 2021; Flamino et al., 2023). While some studies have assessed platform moderation during individual elections (Majó-Vázquez et al., 2021; Chowdhury et al., 2021; Pierri, Luceri, and Ferrara, 2022), few if any have done so across elections.

It is noteworthy that outside of the internet context, political scientists have some normative assumptions for how engagement between different parties in the United States should change between presidential and midterm elections. Particularly, voter enthusiasm is predicted to increase for the party opposing an incumbent president during the midterm election, in this case the Republican party (Campbell, 1991). However, social media metrics and voter turnout have shown to be only weakly correlated (e.g. Cameron, Barrett, and Stewardson, 2016)) and generally underperform traditional polling (Gayo-Avello, 2013; Skoric, Liu, and Jaidka, 2020). Social media metrics are also vulnerable to unique forms of manipulation which may be able to entirely decouple their external signals from voter behaviors (Badawy, Ferrara, and Lerman, 2018; Acker, 2018; Deb et al., 2019). Accordingly, I enter this analysis with little expectation that predictions based on offline voter turnout should carry over to networked publics on Twitter.

5.5 Model

For the purpose of this work, I define the size of a networked public, P, as the sum total of interactions produced by a community of affiliated users over a given time period and on a specific platform. I propose that this interaction total is the product of three factors represented by five variables, displayed in the following formula:

$$P = (N_c * A_c) * (R_c) * (N_a * A_a)$$

The leftmost group of two terms represent the activities of creators, marked by the subscript c. The number of creators is represented by N_c, the rate or activity level at which they produce messages is represented by A_c. The middle term, R_c, represents the intervening force of the medium, or in this case platform, and how this platforms moderates the number of messages which are broadcasted to creators' potential audiences. The rightmost group of two terms represent the activities of audiences within a public, represented by subscript a. Similarly to the creators, N_a represents the number of potential recipients of a message, while A_a represents the sharing rate of these audiences, or their propensity to reshare information holding constant the availability of that information.

Each of these five variables represent *distributions*, rather than scalar values, as individual information creators and audience members will differ from each other in their reach and activity levels. These groups are also slightly overlapping in that creators often also play the roles of audiences for other creators.

I propose that changes in the size, or rather total output, of a networked public can be represented profitably by changes to these five variables. My findings test several hypotheses as to how these variables may have been affected from 2020 to 2022 to explain a primary finding: a decrease in

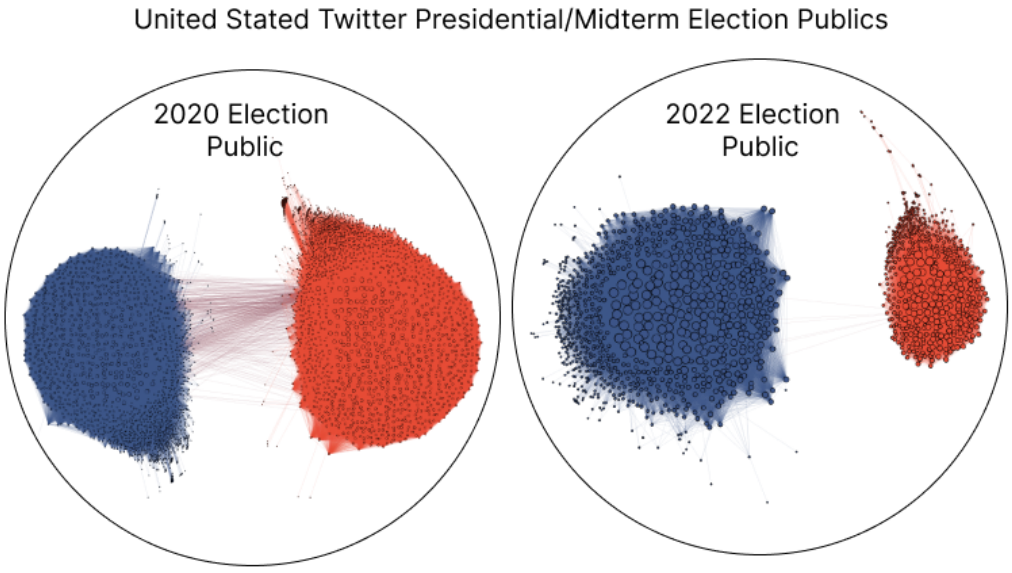

FIGURE 5.1: Visualizations of influencer networks used to classify partisan users in 2020 and 2022. Both networks are created from Twitter posts using the words "vote," "ballot", "election," and their derivatives. Partisanship is determined by users' interaction with right-leaning influencers, colored red, and left-leaning influencers, colored blue. Node size is determined by the influencer's centrality in the network.

communication output as represented by the Twitter retweet (P) in right-leaning election publics, relative to left-leaning publics.

5.6 Datasets

5.6.1 Retweet Datasets

There are two primary datasets for measuring tweet and retweet volume in this study. The first, labeled *Election2020*, spans from September 1, 2020 to November 18, 2020. The second, labeled *Election2022*, spans from the same date range in 2022. The latter date range was chosen because many election-related accounts suspended after the 2020 election were unsuspended on November 18, 2022, creating a potential discontinuity in 2022 election publics (Duffy, 2022).

Both datasets are drawn from larger datasets comprised of tweets collected via the Twitter Streaming API. These larger datasets collected tweets related to a variety of election-related terms, and were updated regularly during the time period of this study to add new terms pertaining to emergent, election-related events. In the present study, I take subsets of these larger dataset that only include the terms, [*'vote'*, *'votes'*, *'ballot'*, *'voted'*, *'voter'*, *'election'*, *'ballots'*, *'voting'*, *'voters'*], all of which were specifically requested from the API in both datasets. Character strings with substrings including these terms are not included, e.g. "#election" and "voterfraud" are not included. *Election2020* totals 408M retweets of 4.1M unique users, while *Election2022* totals 200M retweets across 2.2M unique users.

This analysis requires separating both datasets into right-leaning and left-leaning users respectively. Intuitively, users are classified as either right- and left-leaning if they primarily retweet either popular right- or left-leaning accounts during their respective elections (Figure 5.1). To identify these popular partisan accounts, I construct a coengagement network for both election datasets, where each node is a user and two users are connected if they are both retweeted together at least twice by at least 500 users (Beers et al., 2023a). The results of both graphs are then clustered using the Louvain algorithm (Blondel et al., 2008). Left- and right-leaning clusters are identified as whichever cluster contains the accounts for presidential candidates Joe Biden and Donald Trump in 2020, and then-president Joe Biden and popular 2022 Republican congressional candidate Kari Lake. Other clusters are discarded, and remove nodes who have at least 10% of their connections between clusters. A counting system is then used on every use who made a retweet in these datasets: If 80+% of their retweets of popular users identified in the previous step were right-leaning, those users would be classified as *right-leaning*, and vice versa for *left-leaning* users. If they did not meet the threshold for either party, they were marked as *bipartisan*. If they did not retweet any popular user identified in the previous step, they were marked as *non-partisan*.

5.6.2 User Sample Datasets

This analysis of tweet volume is supplemented with derived datasets of user samples with respect to different user attributes. Particularly, as described below, I generate user samples to analyze differential user suspensions, user search shadowbanning, and users' posting of non-election related content in time periods closely following the 2022 election.

The generation of these user samples was time-sensitive, as new ownership of Twitter had stated both their intent to reverse existing suspension and shadowban status for popular United States users, and their to restrict API access to user data shortly after the election (Thompson, 2022). Because of this, user samples are generated with a preliminary version of the retweet datasets with some differences in date range and term collection, but with otherwise identical methods of e.g. partisan assignment. Specifically, the date range for the datasets used to generate user samples extends to December 1, 2020/2022, which may include tweets from users' Twitter had reinstated after November 18, 2020. Furthermore, the search term criteria for the 2022 dataset is narrower in these preliminary datasets, leading to a smaller though highly overlapping dataset with the primary 2022 dataset used for calculating post volumes above. All stated conclusions hold if either datasets had been used as the primary dataset for analysis, and additional details of the differences between these datasets is provided in the supplement.

Suspension Data

I used the Twitter API to collect the suspension and account status of all users in the 2020 user sample dataset who had not been retweeted in the 2022 dataset, assuming the remainder was not suspended. This data collection took place from January 14th to January 20th, 2023. This is not a perfect representation of which users had been suspended during the 2022 midterms, as Twitter had began reversing suspensions particularly in right-leaning communities due to a change of ownership in November of 2022 (Thompson, 2022). To mitigate this effect, suspension status is manually checked for the top 1,000 most retweeted accounts in 2020 that were not suspended and

yet never retweeted in 2022 before November 18, 2022, when Twitter began to reverse prominent user suspensions. I used either archives of their profile on the Internet Archive or references to their suspension in external media during the dataset time period to infer the status of their account during the 2022 election.

Due to rate-limiting constraints and changes to the Twitter API's access rules preventing additional data collection, suspension status could not be collected for all users. I prioritized users who were first A) the most retweeted and then B) retweeting the most. Accordingly, I was able to acquire the suspension status for any user which was retweeted more than 200 times, i.e. 40% of all retweeted users in 2020, responsible for 99% of all incoming 2020 retweets. I was able to acquire the suspension status for the top 16% of retweeting users, responsible for 68% of all outgoing retweets in 2020.

Shadowbanning Data

Colloquially, shadowbanning has meant methods taken by a platform to decrease the reach of their users without those users' knowledge. By definition, external researchers cannot know the full extent and variety of shadowbanning on any given platform, but there are specific tests one can run to verify that different users are promoted in different ways. One of the simplest manifestations of shadowbanning on Twitter is removal of a users' posts from platform search results (Jaidka, Mukerjee, and Lelkes, 2023). During the 2022 election, Twitter search results had two sections: "Top," which features popular posts relating to a given search, and "Latest," which features the most recent posts relating to a given search. Some users will have their posts removed from the "Top" search results, possibly decreasing their reach. We can detect these users by crafting search queries which only return posts from a certain user ("from:@user"), and then check if their tweets still do not show up in the "Top" search results panel.

To do this, from January 5, 2023 to January 26, 2023 I systematically visit the search result pages for the top 1,000 most retweeted right-leaning and left-leaning users in the 2022 user sample dataset,

and check to see if their posts are returned in the “Top" search bar. The status of some users could not be determined because they were suspended, had set their accounts to private, or had changed their usernames at the time of collection. Furthermore, post-election changes in Twitter moderation strategy supposedly opposing shadowbans in late 2022 means that some users may have stopped being shadowbanned after the 2022 election but before this data collection. I thus caution that these shadowbanning evaluations likely underestimate the extent of shadowbanning during the 2022 election. I also stress that this measure is only one way that Twitter accounts can be shadowbanned, and that some accounts may be facing more (or less) severe sanctions than can be measured.

Full Timeline Data

To understand user behavior when *not* tweeting about election-related topics, I used Twitter’s V2 API from February 7-8th, 2023 to collect additional Twitter activity of a sample of both the most retweeted and most retweeting users in 2020 and 2022. For the most retweeted, I sample timelines from the top 2,000 most retweeted right- and left-leaning users in both the 2020 and 2022 user sample datasets. For the most retweeted, I randomly sample 2,000 right- and left-leaning users who were in the top 50% of retweeting rate across their election periods in 2020 and 2022. For retweeted users, I only collect original posts, excluding retweets and replies. For retweeting users, I only collect retweets. The timelines of some users could not be collected from the V2 API because their accounts were deleted, they had deleted their past posts, or they had been suspended. For the top retweeted users, I able to include 583 right-leaning accounts and 906 left-leaning accounts in 2020, and 891 right-leaning and 942 left-leaning accounts in 2022. For the high-retweeting rate sample, this dataset included 673 right-leaning accounts and 800 left-leaning accounts in 2020, and 931 right-leaning and 911 left-leaning accounts in 2022.

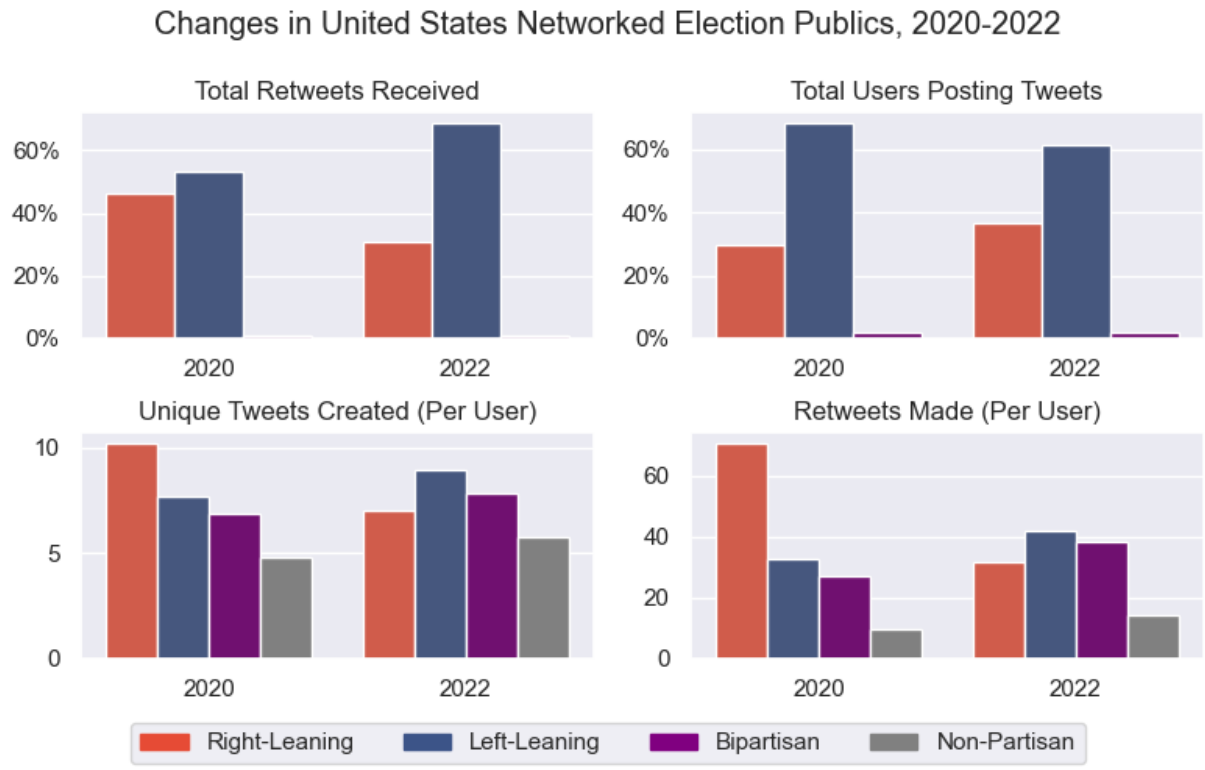

FIGURE 5.2: Four bar graphs describing differences between different partisan groups in the 2020 and 2022 election datasets. The relative number of election-related retweets received by right-leaning users decreased significantly (top-left), but the relative number of right-leaning users making tweets increased slightly (top-right). The relative number of right-leaning tweets and retweets made using election-related terms in 2022 decreased from 2020 (bottom-right and bottom-left).

5.7 Results

In this section I share several hypotheses related to changes in networked election publics between 2020 and 2022, detail my method for testing each hypothesis, and report my findings.

5.7.1 Total Communicative Output (*P*):

I first test my primary hypothesis:

H1: *Retweets of right-leaning users' election-related tweets decreased relative to the left from 2020 to 2022.*

The 2020 presidential election was a much more popular topic than the 2022 congressional elections in these datasets. Election-related retweets dropped from 408M retweets in 2020 to 200M in 2022, a 51% decrease. Within this decrease, however, this hypothesis is supported: right-leaning retweet totals decreased steeply in 2022 relative to left.

In 2020, right-leaning users received 134M retweets (51% of partisan retweets) compared to 155M retweets to left-leaning users (48%). In 2022, the proportion of right-leaning retweeting decreased markedly, with 23M retweets going to right-leaning users (31%) and 59M retweets going to left-leaning users (68%, Figure 5.2). In both years, bipartisan users received less than 1% of all partisan retweets.

5.7.2 Number of Creators (N_c)

To understand this decline in the output of right-leaning publics, the following hypothesis about right-leaning creators is tested:

H2: *The number of right-leaning users posting election-related tweets decreased relative to the left from 2020 to 2022*.

This hypothesis is not supported. From 2020 to 2022, the number of partisan tweet creators overall dropped 48% from 4.1M users to 2.1M users. However, the share of right-leaning creators rose (Figure 5.2). In 2020, there were 521K unique right-leaning users making tweets (30%) compared to 1M unique left-leaning users (68%). In 2022, there were 177K right-leaning users (36%), 298K left-leaning users (61%). In both years, bipartisan users represented 2% of all partisan users. I then measure the effects that contribute to this change in the number of right- and left-leaning creators: *attrition*, *entry*, and *conversion*.

Creator Attrition

Far fewer users chose to tweet about the 2022 midterm elections than the 2020 midterm elections. Both right- and left-leaning publics faced heavy tweet creator attrition, but to slightly different extents. Only 15% of right-leaning creators active in 2020 remained in 2022, compared to 22% of left-leaning creators. Non-partisan users had higher attrition, with only 10% of their 2020 user-base active in 2022. Attrition was much lower for the most popular accounts. Of the 1,000 most retweeted producers in the 2020 dataset, 57% of right-leaning creator remained in the 2022 dataset, and 88% of left-leaning accounts.

There are three reasons a 2020 content producer could be absent in 2022: non-participation, account deletion, or suspension by the platform. Examining the subset of highly-active accounts there is suspension data for, I find that right-leaning attrition was caused much more by platform suspension and account deletion than in other groups (Figure 5.3). Among these right-leaning users present in 2020 and absent in 2022, 35% left because they were suspended, 22% had deleted their accounts, and the remaining 43% of users otherwise did not participate in election discourse. Among absent left-leaning accounts, only 9% of departing users had been suspended, 13% of users deleted their accounts, and 78% of users otherwise did not participate. Non-partisan users had similar levels of suspension and deletion to left-leaning users, with 8% of users suspended, 13% of users deleted, and 78% of accounts not participating.

The suspension differential between right and left is higher among the top creators. Of the top 1,000 most retweeted right-leaning creators from the 2020 sample who did not return in 2022, 79% were suspended, while only 12% of corresponding left-leaning creators were. Overall, 34% of all right-leaning retweets in the 2020 data were given to users who were suspended in 2022, compared to only 3% of retweets for left-leaning creators and 8% for non-partisan creators.

Creators Entry

There are two reasons a tweet creator may newly enter a networked public: they only recently created a Twitter account, or they existed but did not participate in prior publics. Right-leaning communities had higher entry overall than left-leaning communities: 67% of right-leaning creators in 2022 had not participated in 2020, compared to only 46% of left-leaning users. This higher level of entry offset the higher level of suspension and deletion, resulting in the overall gain that right-leaning creators saw in 2022.

Right-leaning creator entry was driven by newly created Twitter accounts, rather than newly-engaged users. I consider users created after the ending date of the 2020 dataset, December 1, 2020, to be new users in the context of the 2022 election. I find that 48% of new right-leaning users in 2022 had created their accounts after the 2020 election, compared to only 23% of new left-leaning users. Left-leaning creators had a similar percentage of new users as non-partisan creators (28%), again suggesting an effect specific to right-wing users.

Creator Conversion

Creators can also switch partisan affiliations, from right to left and left to right and vice versa. In practice, however, this is extremely rare. From 2020 to 2022, 1,000 users switched from right to left, while 2,000 users switched from left to right. These conversions accounted for less than 0.2% of the their overall userbases in 2022.

5.7.3 Number of Audience Members (N_a)

If the number of right-leaning creators did not decrease, it may still be that the number of right-leaning audience members to receive their messages decreased. Twitter, and indeed no popular social media platform, shares information about which users see which content, and so we cannot directly know the size of a potential audience for a given public. Furthermore, many users

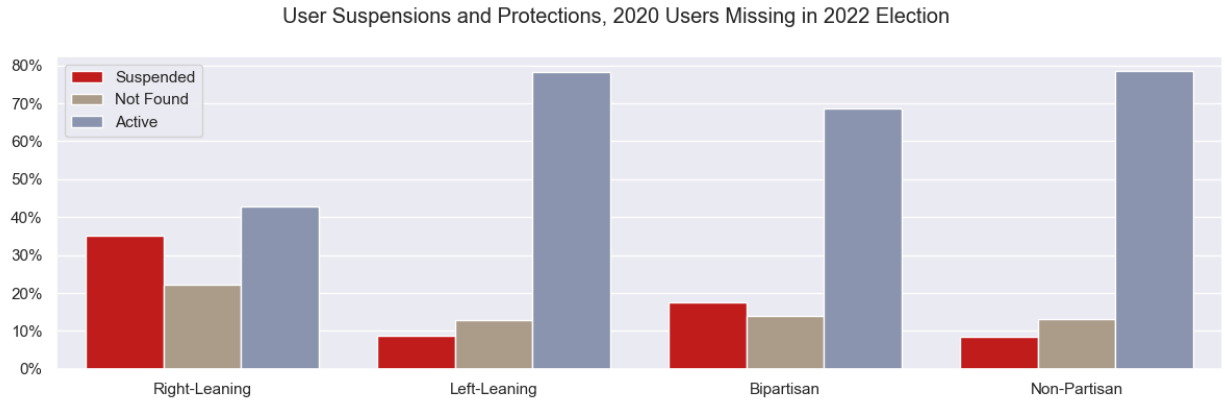

FIGURE 5.3: A bar graph comparing relative rates of suspension, account deletion/removal, and continuing activity in users who were active in the 2020 election, but inactive in the 2022 election.

are exposed to content from creators they would not promote, such as a right-leaning user being exposed to a left-leaning creator. Accordingly, my goal is to measure the audience of users who would *potentially* engage with a creator's content, rather than the total population of users, however disinterested or fundamentally opposed, may be shown that content.

To measure this quantity, I measure the size of an audience as the unique number of users who had retweeted an election-related tweet at least once. These users demonstrate interest in networked election publics, but may be more or less active depending on the year. As a hypothesis:

H3: *The number of unique right-leaning users retweeting election-related posts decreased from 2022 to 2020 compared to left-leaning users..*

This hypothesis is rejected, on largely the same grounds that the right-leaning creator hypothesis was rejected. As with creators, the number of unique retweeting audience members had decreased dramatically from 2022 to 2022, by 44%. And as with creators, the relative share of right-leaning audience members grew from 2020 to 2022. In 2020, 28% of unique partisan retweeters were right-leaning, compared to 70% left-leaning and 2% bipartisan. By 2022, this right-leaning share had

increased to 37%, compared to 61% left-leaning and 2% bipartisan. Thus, while the absolute number of messages retweeted by right-leaning audiences plummeted in 2022, the relative number of right-leaning creators and audience members actually *increased* during this time period.

Attrition, Entry, and Conversion of Audience Members

Trends found in audience members for account attrition, entry, and conversion parallel those found for creators. Left-leaning audience members have higher retention (28%) than right-leaning audience members (21%), as with creators. Of the sample for which there is suspension data, right-leaning audience members have higher levels of account suspensions and deletions (24%) than left-leaning (4%) or non-partisan users (4%). The right again compensates for this higher audience attrition by having more users enter the conversation than the left (64% vs. 38%). This higher audience entry is again driven by a large number of new right-leaning accounts created since the 2020 election, compared to left-leaning accounts (63% vs. 46%).

5.7.4 Reach of Creators (R_c)

Follower Counts

I now assess the reach of creators' content to their audiences, mediated by their platform. On private social media platforms, reach is almost always mediated by non-transparent algorithms which can change at any moment. This prevents external researchers from ever precisely knowing which content is shown to who. However, the platform provides proxy metrics that would seem to be correlated with reach. The most visible metric is users' displayed "follower" counts: the more accounts follow an account, the greater reach its content is likely to have. Accordingly, I first test:

H4a: *Right-leaning creators' follower counts grew more slowly in 2022 relative to left-leaning creators.*.

This hypothesis is supported. In both years, median right-leaning follower counts are much lower than median left-leaning follower counts, which is unsurprising given their smaller overall audience size. However, right-leaning creators' average follower totals grew slightly more slowly from 2020 to 2022 than left-leaning creators. Specifically, right-leaning accounts averaged 53% of left-leaning accounts' follower totals in 2020, but only 51% in 2022.

Shadowbanning

The "following" mechanism is a highly visible platform affordance that mediates the reach of messages from creators to audiences. However, most platform affordances governing reach are largely hidden from users and researchers. One example is investigated here: the practice of *shadowbanning*. To better understand the reach of content producers' posts, the following hypothesis:

H4b: *Popular right-leaning users were hidden from platform search results more often than popular left-leaning users in 2022..*

I find that 108/929 (11.6%) of the top 1000 most retweeted right-leaning accounts data could be retrieved for were search result shadowbanned, compared to 15/957 (1.6%) of the top 1000 left-leaning accounts data could be retrieved for. Thus, in the month after this data was collected, search result shadowbanning was rare, but highly-retweeted right-leaning accounts were shadowbanned at higher rates than left-leaning.

5.7.5 Creator Activity (A_c)

Previous findings indicated that the number of right-leaning creators rose relative to the left in 2022, but it is plausible the average posting rate of these creators could have correspondingly declined. I now test the following hypothesis to understand how active these creators were:

H5: *The number of posts created by right-leaning creators decreased relative to the left from 2020 to 2022.*

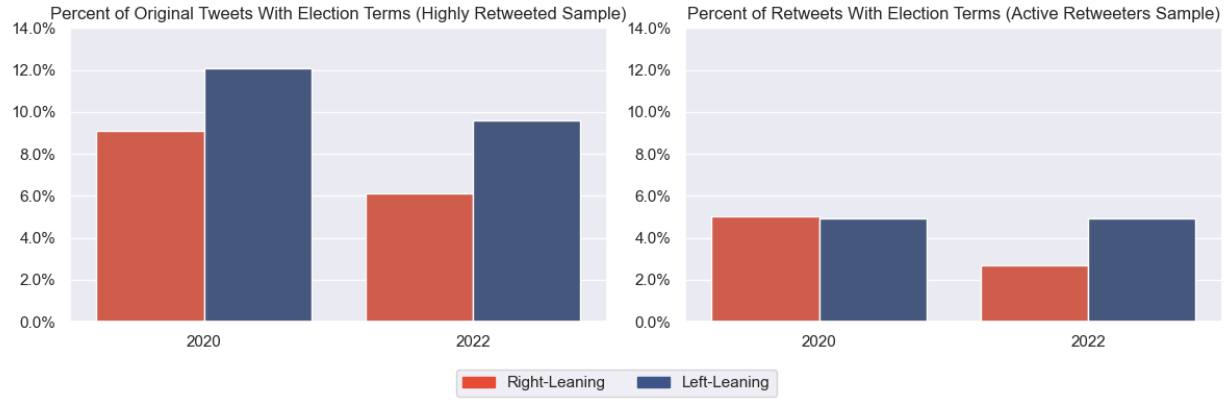

FIGURE 5.4: Each of these bar charts represents a sample of the full timelines of highly-active tweet creators (left) and retweeters (right) during the election. Right-leaning audiences retweeted content using election-related terms less often than left-leaning audiences in 2022, compared to 2020.

This hypothesis is supported; though the relative number of right-leaning creators slightly increased, the number of posts they made decreased (Figure 5.2). In 2020, right-leaning users produced 4.9M election-related tweets (36%) compared to 8.4M election-related tweets by left-leaning users (62%). By 2022, this share of right-leaning tweets had decreased slightly, with right-leaning users creating 1.2M tweets (31%) compared to 2.7M left-leaning tweets (67%). In both years, bipartisan users created 2% of all election-related tweets.

As noted in prior findings, there are far fewer right-leaning creators than left-leaning creators. In 2020, right-leaning creators needed to significantly outproduce left-leaning creators to achieve a similar output of tweets. Accordingly, the mean right-leaning creator made 10.2 tweets in 2020 compared to 7.7 tweets by left-leaning creators. In 2022, however, this pattern flipped, with the mean right-leaning creator making 7.0 tweets and the left-leaning creator making 9.0. This reversal of posting rates lead to the decrease in overall right-leaning posting in the aggregate.

Activity compared to other topics

Right-leaning creators may have decreased their tweet output on election-related tweets specifically, or have simply decreased their output on all topics (Figure 5.4). To address this question, I analyze the full post timelines of a sample of the most retweeted right- and left-leaning accounts from the 2020 election through the 2022 election. I compare the proportion of their total self-authored posts which include words from the original search query (variations of "vote", "ballot," and "election").

Both right- and left-leaning partisan accounts decreased their overall posting output from 2020 to 2022. Right-leaning accounts decreased on average from 1,303 posts to 1,110 posts (-15%), while left-leaning accounts decreased from 1,523 posts to 1,212 posts (-21%). The share of election-related posts from this total also dropped from 2020 to 2022. In 2020, 10% of right-leaning posts contained election-related terms, while 12% of left-leaning posts did. In 2022, the proportion of posts containing election-related terms dropped to 6% in right-leaning accounts and 10% in left-leaning accounts — a slightly higher drop on the right.

5.7.6 Audience Activity (A_a)

I conclude by analyzing the final factor, the audience rebroadcast rate:

H6a: *Right-leaning audiences retweeted election-related content at a lower rate than left-leaning audiences in 2022, compared to 2020..*

This hypothesis is supported (Figure 5.2). The average retweeting audience member in 2020 on the right made more than 2x the number of retweets than the average left-leaning member (71 tweets right-leaning vs. 33 tweets left). By 2022, that statistic had reversed, with left-leaning audience members retweeting slightly more than right-leaning audience members (32 right-leaning tweets vs. 42 tweets left). Had right-leaning audience members retweeted at similar rates compared

to the left in 2022 as they had in 2020, then the overall partisan balance would have completely reversed in 2022, favoring right-leaning tweets.

As described in the previous finding on creators, right-leaning audiences also simply had a lower supply of election-related tweets to retweet in 2022 than their left-leaning counterparts. To account for this, I also correct for the availability of right-leaning retweets as an alternative metric. In both 2020 and 2022, the average right-leaning audience member retweeted a relatively larger share of the available right-leaning tweets than in left-left leaning publics. However, even controlling for the constrained supply of original election-related tweets, right-leaning audience engagement with election content decreased relative to left in 2022.

Activity compared to other topics

As with tweet creators, I pull a sample of the full historic timelines of a sample of high-retweeting audience members from the Retweet Datasets using the Twitter V2 API. On average, this sample of right- and left-leaning users made almost the same proportion of election-related retweets in 2020, but that right-leaning users decreased their retweets of election-related posts in 2022 (Figure 5.4). In 2020, right-leaning users had about 5.0% of their retweets using election-related term, while left-leaning users had 4.9% of their retweets election-related. Meanwhile, in 2022, right-leaning users in this sample decreased their engagement with election-related content to 2.7% of all their content, while left-leaning users only reduced to 4.1%. In both 2020 and 2022, right-leaning users made more retweets than left-leaning on non-election content, suggesting that this decrease in right-leaning activity was specific to election-related content.

Followback Networks

These findings on audience activity are complicated by the fact that all audience members do not retweet alike. Specifically, previous research have found that certain users are members of retweet spam networks called "followback networks," which retweet at a much higher rate than

other users (Gallagher, 2020; Torres-Lugo, Yang, and Menczer, 2022; Beers et al., 2023a). These groups of users seek to exploit Twitter's norms of reciprocal following by following as many users as possible, expecting to be followed back. They support each other in this effort by retweeting lists of other users' screen names that they know are likely to follow back, flagging them for their fellow followback members. These communities were particularly active in the 2020 United States election, supporting both right- and left-leaning communities, and were marked by much higher rates of retweeting than non-followback communities (Beers et al., 2023a). Torres-Lugo, Yang, and Menczer (2022) observed that pro-Trump followback communities in early 2020 frequently shared low-credibility news and conspiracy theories, and observed structurally similar followback communities among anti-Trump users.

The existence of followback networks stresses an important qualification to audience-related factors in this model: audiences are not homogeneous. Changes in retweet output can reflect changes in one sub-audience, but not another. Accordingly, to measure the extent to which followback networks in particular changed from 2020 to 2022, I test the following hypothesis:

H6b: *Right-leaning followback networks decreased in size and activity levels in 2022 relative to left-leaning followback networks*.

This hypothesis affirmed on both of its components. To isolate followback accounts in these datasets, two criteria are applied. The first is that the number of accounts a user is following must be at least 90% of their current follower count. The number of users an account is allowed to follow is limited by their current number of followers, meaning that followback accounts will always be following somewhere near their current follower count. The second criteria is that the user must be following at least 5,000 users, as this is the boundary where the first criteria begins to apply.

Using these criteria, in 2020, there were 32K right-leaning followback accounts and 32K left-leaning followback accounts, comprising 2% and 1% of all partisan users respectively. As expected, these accounts on average retweeted much more than non-followback accounts: 565 retweets during the 2020 election season for right-leaning followbacks, and 288 retweets for left-leaning followbacks. Overall, right-leaning followbacks produced 18M retweets over the 2020 dataset for 13% of all right-leaning retweets, compared to 9M retweets created by left-leaning followbacks for 6% of their retweets. Clearly, in 2020, followback accounts accounted for a disproportionate number of overall retweets relative to their small numbers. These right-leaning followback accounts not only comprised a bigger portion of their community, but were also almost twice as active as left-leaning followback accounts.

In 2022 election discourse, the right-left distribution of followback accounts entirely reversed. This time, there were more than twice as many left-wing followbacks as right-leaning followbacks (26K vs. 11K). Left-wing followbacks were *more* active than right-leaning followbacks in 2022 (231 vs. 135 average retweets), although in general all partisan followbacks were less active in 2022 than in 2020. Left-wing followbacks accounted for 11% of all left-wing retweets, compared to only 6% for right-leaning followbacks. In effect, left-wing followback activity as a proportion of all retweeting activity increased in 2022, while right-leaning followback activity as a similar proportion steeply decreased.

The lack of right-leaning followback activity can be related to differential suspension rates between right-leaning and left-leaning accounts after 2020. As previously described, right-leaning suspensions rates for audiences and creators in the retweet datasets were more than 3x higher than left-leaning suspension rates after 2022. This trend is even more acute for followback spam networks. In a sample of the 8.5K most active left-leaning followback accounts in 2020, they were suspended in similar proportions to non-followback accounts at 13%. By contrast, the right-leaning followback suspension rate was higher than the non-followback level of right-leaning suspension, at 52% of all followbacks in 2020 (sample size: 15.6K accounts).

5.8 Discussion and Conclusion

Two datasets of Twitter posts are presented relating to the 2020 and 2022 US elections, and analyzed to understand how right and left-leaning networked publics have changed across five broad metrics. Two somewhat contradictory trends are revealed with respect to right-leaning publics in this study. The first trend suggest a recession of the right-leaning networked election public after their widely-publicized engagement with election fraud conspiracy theories in 2020, while the second suggests tentative growth. First, despite positive indicators, right-leaning posting and audience retweeting activity about elections declined from 2020 to 2022 relative to left-leaning activity, and right-leaning accounts in 2020 missing in 2022 were more likely to be suspended than similar left-leaning accounts. Decreased retweeting activity and account suspension were particularly severe in right-leaning retweet spam ("followback") networks, which were highly active in 2020 (and implicated in the spread of conspiracy theories (Torres-Lugo, Yang, and Menczer, 2022)), relatively dormant in 2022, and moderated more severely than corresponding left-leaning followback networks.

The second trend identified in this study, however, suggests growth in the number of participants in the right-leaning election publics. Relatively more unique right-leaning users participated in 2022 election publics than did left-leaning users. While right-leaning audiences and creators decreased their activity relative to the left during the election, their activity levels were higher than the left on non-election topics. Overall this study portrays a potentially growing right-leaning public that nonetheless failed to mobilize for election discourse at both the creator and audience level, compared to their left-leaning counterparts.

In presenting these empirical findings, I develop a model of measuring publics that flows from creators to audiences via the medium of platforms. We must not make the mistake, however, of believing these factors are independent. Audience interest and message production are often cyclically linked, with increased interest stimulating additional message production and subsequent

increased messaging stimulating audience interest (Feezell, 2018; Barberá et al., 2019). Furthermore, many creators, especially at the elite level, are motivated to produce messages that their platforms will promote, rather than restrict, in order to increase their engagement, maintain audience share, and avoid suspension (Ma, Gui, and Kou, 2023). In the election context, where right-wing publics' denial of election results had resulted in suspension and moderation in 2020, such interplay between creators and platform reach is especially pertinent. The use of this model is thus only a starting point for deeper qualitative and experimental research on the complex incentives for change that actors within networked publics face.

The drawbacks of measuring changes in networked election publics by their number of retweets are clear. As I have noted in the findings on followback networks, the metric is easily manipulated by organized spam networks, who distort the presumed correlation between generalized audience interest and retweet totals. Furthermore, the capricious influence of unknowable industry algorithms constantly changes how platforms mediate messages, further uncoupling retweet totals from generalized audience interest. Such metrics should thus not be hastily applied to external phenomena, such as voter enthusiasm or intent, but rather should be taken as outward reflections (and distortions) of the inner dynamics of a distinct networked public.

Despite this caveat, externally visible metrics like the retweet are the metrics that larger, traditional publics can see, and thus become the grounding for popular media discourse and academic research on the health of online communities. Users themselves assiduously track such metrics, and derive meaning from the perceived relationships between different social media metrics (Guan, 2019). Journalists cite retweet numbers as indicators of the popularity of a given tweet or account, and use "popular" tweets to frame public opinion (McGregor, 2019). While external metrics are artificial and easily manipulated, their high visibility constructs a social reality with real consequences. I encourage future research to both understand and publicly communicate the dynamics that undergird the generation of different social media metrics, and how well they relate to the qualities they are widely perceived to measure.

This study faces a limitation which is common to many social media studies but particularly acute in the present moment: the problem of working with data provided by a private company with no incentives for data transparency. My analysis of suspensions and shadowbanning, for example, uses data which is deliberately obscured from users and researchers by Twitter. A consequence of this lack of transparency is that Twitter can arbitrarily reverse suspensions and shadowbanning, leaving external auditors little evidence that they had occurred in the first place. As decision makers at private social media companies increasingly restrict access to their data, external auditors, whether part of academia, the government, or civil society, will likely have to face increasing compromises in what can be known (Jingnan, 2023). Studies such as this may well represent a near-term high watermark of public data provenance about our networked publics, and future researchers will need to use creativity to investigate longitudinal change in these publics.

5.9 Acknowledgements

Funding for this work has come from the University of Washington's Center for an Informed Public, the John S. and James L. Knight Foundation (G-2019-58788), Craig Newmark Philanthropies, the Election Trust Initiative, and the National Science Foundation (grant #1749815 and grant #2120496).

5.10 Appendix: Construction of Datasets

There are four primary datasets used in this paper. The first two, the Retweet Datasets, provide the raw retweet data for 2020 and 2022, and is responsible for all retweet-volume statistics in this dataset. The second two, the User Sample Datasets, provide the basis for generating appropriate users samples for suspension, shadowbanning, and full timelines analysis in either 2020 or 2022.

The User Sample Datasets were time-sensitive to collect. The owner of X, formerly Twitter, had made statements that they were reversing suspension and shadowbanning designations in the

months after the 2022 election. Additionally, the API access upon which this study relied was cut off in the months after the 2022 election, making later collection of suspension and timelines impossible. Accordingly, these User Sample Datasets were collected with a provisional version of the 2020 and 2022 datasets which would later form the basis of the Retweet Datasets.

I first describe the basic procedure for collection of all four datasets. All four datasets are collected using the Twitter V1 Streaming API. This API streams tweets mentioning certain terms over a given time period. In this case, this dataset was set to collect terms relating to the United States elections in 2020 and 2022. The specific terms collected in these datasets were updated over time, as different developments happened in both elections, and paid particular attention to terms relating to the spread of election-related misinformation. As described in the full paper, datasets with these larger collections of terms are then subsetted to only include [*'vote'*, *'votes'*, *'ballot'*, *'voted'*, *'voter'*, *'election'*, *'ballots'*, *'voting'*, *'voters'*]. These terms are intentionally collected with word boundaries, meaning, for example, the hashtag #voterfraud would not be included among them. My goal in using word boundaries was to avoid viral terms which might have a specific partisan bias.

While both the Retweet Datasets and the User Sample Datasets use this same core method, the Retweet Datasets improve upon the coverage of the User Sample Datasets in several ways. This makes them more appropriate for analyses of retweet volume, but make them differ in some respects from the User Sample Datasets.

First, I restricted the date range in the Retweet Datasets to end on November 18, 2020/2022 instead of December 1, 2020/2022, to align with when suspension reversals of United States election influencers were purported to be happening. My goal with this change was to not be measuring two different regimes of particularly right-leaning influencer activity in 2022.

Second, the 2022 User Sample Dataset originally pulled from a smaller and slightly different pool

Dataset	2020 Retweets	2020 Retweet Overlap	2020 Retweet Overlap Percent	2022 Retweets	2022 Retweet Overlap	2022 Retweet Overlap Percent	Start Date	End Date
Retweet Dataset	407783003	407303485	99.88%	199507224	55994407	28.07%	Sept 1	Nov 18
User Sample Dataset	483339828	407303485	84.27%	66533273	55994407	84.16%	Sept 1	Dec 1

TABLE 5.1: Differences between datasets used in this paper. "Overlap" refers to the number of tweets found in both the corresponding-year Retweet and User Sample datasets.

of terms streamed from the Twitter V1 API, reducing both the size and variety of its content. Particularly, some of the terms extracted, including variations of "vote" and "ballot' were not specifically requested from the API in this first version of the dataset. This is corrected in the 2022 Retweet Dataset, resulting in a much larger and likely more representative dataset for analysis, which specifically requests from the API all of the terms that are filtered for in this study.

Third, the methods for detecting word boundaries between the two datasets were slightly changed, resulting in minute differences in tweets collected.

Differences between the Retweet and User Sample Datasets are covered in 5.1. The datasets for 2020 are nearly identical besides the change in date range. The User Sample Dataset for 2022 is approximately a third the size of the Retweet Dataset, with heavy overlap in retweets in the smaller dataset. All the conclusions in this paper hold if the Retweet Dataset is substituted for the User Sample Dataset. User samples derived from the Retweet Dataset instead of the User Sample Dataset have high overlap; for example, the Top 1,000 most retweeted users in right-leaning and left-leaning clusters have 80+% overlap in both partisan groupings.

The following are lists of terms specifically requested from the API for all four datasets.

2020 Retweet and User Sample Dataset API Collection Terms

vote, voter, votes, voted, ballot, ballots, election, rigged, rigging, rig, riggedelection, voting, Destino2020, el debate, VotoLatino, Decision2020, debate final, elecciones2020, eleccion presidencial, georgia, sharpiegate, maricopa, stopthesteal, ballotgate, recount, nevada, stopthesteal, sharpiegate, ballotgate, carolina, wisconsin, florida, break the windows, waited line, disenfranchise, thestealison, 2020winner, longlines, sorosownedusps, stopthecheat, long lines, foreign interference, redmirage, 2020election, silentmajority, disenfranchised, Trump2020landslide, thebigsteal, yourpresident, stopthe cheat, electionday, longline, waiting line, Soros owned usps, mirage, long line, malfunction, provisional, notmypresident, breakthewindows, 2020vote, vote2020, detroit, steal, 2020presidentialelection, election2020, rigger, benford, blue mirage, VoteBidenHarrisToSaveAmerica, sarasota, maricopa, michigan, bluetsunami, blue tsunami, kent, redtsunami, redmirage, arizona, red mirage, minnesota, hampshire, broward, bluemirage, red tsunami, convoys, militia, foreigninterference, armed, oathkeepers, caravan, shooter, convoy, riot, minutemen, brooks brothers, brooksbrothersriot, proudboys, 4moreyears, parade, shooting, national guard, nationalguard, boogaloos, riots, intimidating, antifa, intimidation, proud boys, intimidate, open carry, blm, boogaloo, brooksbrothers, protesters, fourmoreyears, caravans, opencarry, protests, parades, biden, biden2020, joebiden, kampala, harris, kamalaharris, kamala, jojorgensen, hunter, berniesanders, elizabethwarren, burisma, sanders, warren, kanye, kanyewest, DNC, bernie, jorgensen

2022 User Sample Dataset API Collection Terms

midterms2022, midterm, elections, elections2022, midterms, election2022, election, desantis, midterms2022, midterm, elections, elections2022, midterms, election2022, integrity, RNC, USPS, conservative, GOP, postal, GOPers, DNC, conservatives, liberals, liberal, republicans, republican, dem, democrats, dems, democrat, trump, biden

2022 Retweet Dataset API Collection Terms

midterms2022, midterm, elections, elections2022, midterms, election2022, election, desantis, midterms2022, midterm, elections, elections2022, midterms, election2022, decertification, touchscreen, touchscreens, BMD, vulnerability, tabulators, vulnerabilities, EVM, tampered, tabulator, fraud, suppression, adjudication, electionfraud, tamper, chain of custody, polling, pollbooks, tampering, pollworker, recount, precinct, forensic, EVMs, intimidation, cisa, overvote, pollwatchers, electioneering, votesuppression, fraudulent, polls, hand count, votersuppression, pollbook, absentee, electors, decertify, BMDs, rigged, voterfraud, risk-limiting audit, poll, epoll, undervote, subversion, ballot, ballots, vote, voting, voted, voters, voter, votes, tallies, rolls, fulton, pinal, racine, philadelphia, detroit, maricopa, michigan, pima, cochise, lancaster, arizona, pennsylvania, chicago, yuma

Chapter 6

Study 3: Negativity in U.S. Influencer Culture

Contributing Authors: Anna Beers[a], Julie Vera[a], Gauri Nayak[a], Trushaa Ramanan[b], Sehrish Daud[b], Kate Starbird[a], Emma S. Spiro[b]

[a]University of Washington, Human Centered Design and Engineering

[b]University of Washington, Information School

6.1 Preamble

This chapter is in preparation for a journal submission, and this is the first instance in which this text has appeared. However, abstracts relating to aspects of this project have been submitted to two conferences as of writing, and an earlier version of this project was presented in (Beers, 2023). This preamble serves as additional contextualization for this study within the context of this dissertation, and personal reflection on my part for where this work has since travelled. I use single-spacing in this preamble to offset it from the main text.

This study was directly inspired by Study 4, and is chronologically the most recent study. Particularly Study 4 consisted of an analysis of negativity in the use of citation, and degradation of certain information sources. I increasingly was seeing influencers as information sources themselves, and so was interested in how influencers might degrade each other. It also contributed to one of the main takeaways from Study 3, which is that information does not exist without its other.

In the context of this dissertation, this study asks questions about what strategies influencers use to boost engagement, and how those strategies may alter political discourse in a community.

The original plan for this study was also to conduct an analysis of negativity with respect to racial and gender groupings, as the current Discussion section describes. On the one hand, the results would likely not be surprising, as there is ample literature using multiple methods that shows that women and marginalized racial groups are consistently disproportionately targeted for negativity. However, I think it would be an interesting beginning to a thread of racial and gender power analysis in my studies of influencers, already in explored in papers such as (Christin and Lu, 2023). I am frequently compelled by the (unsurprising) whiteness and maleness of popular "progressive" media, and understanding mechanisms of racialized negativity and their effects on follower growth might prove a productive avenue to better characterize these racial hierarchies.

This study uses a familiar research model which uses group qualitative coding as a means to train a natural language processing model that can then do downstream tasks. There is some reasonable frustration with this model of research, as many have observed that qualitative coding is inextricably tied to the identity of context of the coders, and thus models trained from those coders will inevitably encode their biases. This is to say nothing of the lack of generalizability to new contexts, which is also certainly present in this data. I agree with both complaints, and in this particular study, believe that our coders should be taken to represent the perspective of our team, rather than "truth," and similarly that our model should not be used without pre-training on new datasets. I think a point for improvement for this paper is to more explicitly advocate for this qualitative ethic to a quantitative method: that coding to create data for an algorithm is a way of heightening a team's specificity, rather than making their work "generalizable."

6.2 Abstract

One method for social media influencers to gain attention and exert influence online is through the application of targeted negativity: harassment, feuding, disagreement, controversy, "drama." In this study, we use two years of engagement data tracking 1,000+ influencers affiliated with three separate ideological communities in U.S. discourse, and develop a classification model to determine when influencers are targeting each other for negative content. We test several hypotheses related to the targeting of negativity, and the benefits of negativity as measured by follower growth and post engagement. We find that influencers more often target negativity towards out-groups and receive more engagement when they do so. The recipients of negativity are unevenly distributed, with some influencers receiving far more negativity than one would expect given their

popularity. Furthermore, different ideological groups use negativity in different ways. We discuss the implications of negativity in influencer communities with respect to norm enforcement and proximity to the "mainstream."

6.3 Introduction

Popular United States media has frequently portrayed social media platforms as places where disagreement, anger, toxicity, hate speech, and other forms of negative content thrive (Duggan, 2017; Feder, 2021; Edsall, 2022; Angwin, 2023). Recent incidents where online discourse has spilled over into offline action has brought increasing attention to the moderation of online communities, and the possibility that toxic behavior online may be threatening democratic norms and institutions (Timberg, Dwoskin, and Albergotti, 2021; Edsall, 2021; *Unsafe* 2024). A recurring character in the drama of online negativity is the social media influencer, who alternately directs their followers to harass certain targets (Marwick, 2021), is targeted themselves for mass shaming in episodes of "cancel culture" (Clark, 2020), or who engages in performative "feuds" with other influencers to stimulate audience engagement (Christin and Lewis, 2021; Lewis and Christin, 2022; Harris, Foxman, and Partin, 2023). The past decade has seen a rise in the comparative power of social media influencers, particularly in political discourse, making them an increasingly indispensable part of how negativity is routed and re-routed in online ecosystems (Robb, 2020; Riedl, Lukito, and Woolley, 2023; Hund, 2023).

In this article, we use two years of Twitter data tracking the total engagement data of 1,848 United States (U.S.) influencers to understand how they use, target, and potentially profit from negativity online. We use network methods to classify these influencers into three different ideological communities, and analyze how negativity is used between members of different communities *and* towards members of one's own community. To understand negativity, we introduce a dataset of 2,059 quote tweets hand-coded for negative content, and train a language model based on these tweets to classify over 1.4M additional quote tweets collectively retweeted over 100M times. To

understand why influencers use negativity, we evaluate a regression model comparing their use of negativity longitudinally and what effect it has on their overall follower growth from 2021-2023.

We test four groups of hypotheses, in addition to performing a wider case study on the targets and topics to which negativity is directed in contemporary U.S. influencer culture. We generally find that negativity is disproportionately directed at out-groups rather than in-groups, and furthermore at a few individuals within them. Negativity is an effective way to increase engagement when targeting out-groups, but less so when targeting in-groups,

6.4 Research Questions

The first group of hypotheses concerns the use of negativity between groups, and whether influencers are incentivized to target out-groups members when compared to in-group members. These hypotheses essentially test the core tenets of social identity theory (Tajfel et al., 1979; Tajfel, 1982), which posits that in-groups use negativity against out-groups (and positivist towards in-groups) to maintain social status.

H1a: *Influencers' quote tweets will be more often negative when directed at out-group members than at in-group members.*
H1b: *Negative quote tweets aimed at out-group members will receive more engagement than those aimed at in-group members.*

The second group concerns the disproportionate targeting within groups. We predict that negativity is not evenly distributed across members of out-groups, even accounting for different levels of popularity among out-groups members. That is, distinctive attributes and behaviors of certain users will attract negativity, supporting elements of the morally-motivated networked harassment hypothesis advanced by Marwick (2021).

H2a: *Some influencers will be disproportionately targeted for negativity by out-groups, even when controlling for their overall popularity..*

H2b: *Some influencers will be disproportionately targeted for negativity by their own in-group, even when controlling for their overall popularity..*

The third group of hypotheses concerns the potential benefits — and drawbacks — of negativity for influencers within an attention economy. On the one hand, some influencers may increase their following by participating in sensationalized "drama" or feuds with other influencers, especially when the negativity they received is from out-groups (Christin and Lewis, 2021; Lewis and Christin, 2022; Harris, Foxman, and Partin, 2023). By contrast, some influencers may lose status from incoming negativity, due to social ostracization from their peers (Marwick, 2021).

H3a: *Influencers' who receive more negative quote tweet engagement from out-groups will see their followers increase more rapidly than those who are targeted less*.
H3b: *Influencers' who receive more negative quote tweet engagement from in-groups will see their followers increase less rapidly than those who are targeted less*.

Finally, we consider the possibility that different online communities will have different propensities to use negativity, and benefit from negativity to different extents. Many have theorized that certain U.S. political cultures, particularly far-right groups represented in this dataset, engage on social media more negatively (e.g. Brady et al., 2019).

H4a: *Some ideological groups will use negativity at higher rates than others*.
H4b: *The rate of outgoing negativity between different ideological groups will differ from each other..*
H4b: *The benefits of sending and receiving negativity will differ between different groups..*

6.5 Literature Review

6.5.1 Social Identity, Negativity Bias, and Shaming

The expected operation of negativity between ideological communities derives from one of the tenets of social identity and categorization theory: that people tend to endorse members of their

own group to increase their status, and more conditionally derogate members of out-groups to do the same (Tajfel et al., 1979; Tajfel, 1982; Brewer, 2001). These principles have been shown to operate on social media and particularly Twitter (Wojcieszak et al., 2022), but perhaps more attention has been given to negativity and out-group bias. Particularly, large studies of social media have repeatedly shown that content which either uses negative (and particularly angry) language, or derogates an out-group, tends to out-perform other types of content, particularly in political contexts (*Partisan Conflict and Congressional Outreach* 2017; Rathje, Van Bavel, and Linden, 2021; Robertson et al., 2023, though see Burton, Cruz, and Hahn, 2021). While these studies have focused on the behavior of mass audiences, the corollary likely holds true for influencers seeking audiences' attention: they will more likely to make posts targeting, and more likely to be negative targeting the out-group than the in-group.

Existing social media research on social identity theory and negativity bias provide an explanation for the prevalence and causes of negative out-group targeting, but they speak less directly to negative in-group targeting. Here a framework of in-group punishment or *shaming* is preferred. The psychology literature has studied in-group punishment as a potentially prosocial measure to facilitate group cooperation (Oliver, 1980; Fehr and Gächter, 2002). However, Klonick, 2015 and others have posited that while in-group punishment may have once had generally prosocial functions, the affordances of the internet have greatly increased the frequency of punishment while greatly reducing its prosocial benefits. Shaming, sometimes called "cancelling," is now increasingly operationalized as something irrational and/or based in prejudice that designers of social systems should strive to minimize (Muir, Roberts, and Sheridan, 2021; Vogels et al., 2021). As we discuss in the following section, however, social media influencers have often productively used shaming and in-group punishment to achieve activist aims and reinforce group norms, complicating the supposed negative valence of shaming (Nakamura, 2015; Clark, 2020).

6.5.2 Influencer Drama, Harassment, and Norm Enforcement/Performance

This study centers on the figure of the political influencer, a rising force in U.S. political media particularly among the young (Robb, 2020; Newman et al., 2023). The study of influencers evolved from the study of celebrities, with the early internet term of microcelebrity gradually evolving into the contemporary term influencer (Marwick and boyd, 2011; Senft, 2013; Abidin, 2018). Regardless of terminology, the concept always represented both a status and a practice: a state of being popular on certain online media platforms, and aspiration to be popular reflected in one's actions on those platforms. As media platforms and advertising markets increasingly monetized the status of being popular online, the act and status of being an influencer increasingly represented the status of seeking and achieving economic security and power (Duffy, 2017; Abidin, 2018; Hund, 2023).

One of the strategies to acquired this online capital is the *performance* of norm enforcement online. Christin and Lewis have described the practice of "platform drama" on social media as a way to garner attention and maintain status in influencer communities (Christin and Lewis, 2021; Lewis and Christin, 2022). Lewis and Christin (2022) describe platform drama, the public spectacle of influencers feuds and the commentary thereupon, as "rituals through which the values of YouTubers get expressed, contested, and reaffirmed" (1651). Perhaps more pervasive than platform drama however is simply the practice of influencers negotiating politics online, increasingly called being a "political influencer" (Lewis, 2018; Riedl, Lukito, and Woolley, 2023). Political influencers make a daily practice of expressing, contesting, and reaffirming political values in order to grow and maintain their audience, and accordingly negativity and "drama" is a crucial weapon in their arsenal (Harris, Foxman, and Partin, 2023).

The amount of "toxicity," i.e. negativity, in online political discourse is often decried as a social problem — an inappropriately high level of emotionality which degrades the rationality needed for productive debates about policy (Jamieson et al., 2017; Edsall, 2022). There is a rich literature

on online influencers' use of negativity to demonize and harass individuals and communities, particularly women and historically marginalized communities (Gray, Buyukozturk, and Hill, 2017; Massanari, 2017; Cross, 2019). However, it should be noted some of the greatest triumphs of digital culture have been achieved through conspicuous negativity. Many activists, such as those in the struggle against anti-Black racism in the U.S. (Freelon, McIlwain, and Clark, 2016; Jackson, Bailey, and Welles, 2020; Steele, 2021), use "call-outs" and negativity to fuel and achieve tangible results from social movements. In fact, viewed in this light, the regular use of negativity, alternatively frame as *accountability*, can be seen as an (usually unpaid) service in online communities, enervating solidarity while chastising and reforming users whose antisocial behavior threatens marginalized users and groups (Lorde, 1997; Nakamura, 2015; Clark, 2020). Accordingly, in this study, we do not perceive negativity as a problem to be minimized, but rather a rhetorical strategy that can be used contextually for both pro- and antisocial aims. We also favor the term negativity (compared to other contemporaries such as incivility (Jamieson et al., 2017), impoliteness (Graham and Hardaker, 2017), or toxicity (Fortuna, Soler, and Wanner, 2020)) for its lack of valence and flexibility when used for different pro- and anti-social aims.

6.5.3 Quote Tweets, Stance Detection, and Negative Networks

There exists a substantial literature on the use of the quote tweet, most recently reviewed by (Bastian, 2023). Bastian found that prior research tends to find that quote tweets generally receive more attention (e.g. via incoming retweets) and influence, variously defined, than non-quote tweets (Garimella, Weber, and De Choudhury, 2016; Zheng et al., 2020; Matalon et al., 2021).

Given that quote tweets were almost immediately identified as a way to express opinions vis a vis their targeted tweets, many have been interested in analyzing and automatically classifying the *stance* of tweets AlDayel and Magdy, 2021. Several of these efforts have produced publicly-available datasets manually annotating tweets for stance (Mohammad et al., 2016; Conforti et al.,

2020), and some particularly oriented towards quote tweets (Villa-Cox et al., 2020). However, almost all of the stance detection dataset to date have noted a user's stance with respect to the topic or entity being discussed in the quote tweet, rather than the author of the quote tweet. A notable exception is (Wojcieszak et al., 2022), which does code and develop a model for stance towards the quoted tweet's author, but has not yet made their manually-coded labels available. Wojcieszak et al. (2022) analyzed right- and left-leaning U.S. Twitter users, and found that both groups were more likely to attach negative quote tweets to out-group members than in-group members, some types of accounts (e.g. politicians) were more likely to receive negative quote tweets, and that certain topics were likely to attract more negative quote tweets from certain groups (e.g. immigration policy). While many models have been developed for classifying stance in tweets and quote tweets across a variety of topics, contexts, and actors, model accuracy remains low, rarely reaching an F-score higher than 0.7.

The process of mapping negative associations between individuals has also been approached from the perspective of network science. Usually, the application of social networks to social media data has been to identify certain users or entities as nodes, and social interactions or relations between them as edges. The dominant paradigm in social network analysis has assumed an *unsigned* network where edges between nodes suggest at minimum the possibility of similarity, attraction, or coordination between those two nodes. This dominance likely stems from the more sensitive nature of collecting data about dislike relationships (Asher and Dodge, 1986), and particularly the paucity of publicly-available dislike interaction data on social media platforms, in comparison to positively- or neutrally-valenced interactions (Gray, 2021). To the extent signed network analyses have been done on social media platforms, they have been mostly limited to lesser-used platforms where negatively-valenced interaction data is made available (Kunegis, Lommatzsch, and Bauckhage, 2009; Leskovec, Huttenlocher, and Kleinberg, 2010).

We use two concepts from the existing literature on signed networks: structural balance theory

and notions of signed centrality. Structural balance theory (Cartwright and Harary, 1956) suggests that members of signed social networks generally tend to form connections according to various principles of relationship *balance* over the long-term. Particularly relevant to our paper are the notions that it is most likely that the enemy of an friend is an enemy, or conversely that a friend of a friend is likely also a friend. Accepting this principle of signed networks lets us generalize from influencers to their audiences, suggesting, for example, that audiences form negative opinions about the influencers disliked by their co-partisan influencers. We also borrow notions of signed centrality from the previous literature, particularly the straightforward metrics of *popularity* of nodes as a ratio of the positive incoming interactions a node receives over all the incoming interactions a node receives Zolfaghar and Aghaie, 2010; Tang et al., 2017. We use the inverse in this paper, a dislike score, to mean the ratio of negative interactions an influencer receives out of all of their interactions.

6.6 Datasets and Methods

6.6.1 *Two Years of Influence* Dataset

The dataset underlying these results is titled *Two Years of Influence* (*TYoI*), and represents the entire post history of 1,848 popular Twitter users from March 2021 to March 2023, including all posts of other users who retweeted them, quote tweeted them, replied to them, or retweeted quotes or replies to them. This dataset total 34M original tweets by influencers and 2.7B tweets retweeting, quoting, or replying to content by these influencers. We exclude two months for analysis here, March 2021 due to incomplete data, and November 2022 due to a data collection error, for a total of 22 months analyzed from April 15, 2021 to March 15, 2021.

The influencers tracked in this dataset were subsetted from a larger set of 5,218 influencers gathered from a large Twitter dataset collected before and after the 2020 United States election (Tollefson, 2021). This larger set of influencers was identified through the creation of a co-engagement

graph, which is a network method for identifying communities of popular social media users who share large audiences of other users (Beers et al., 2023a). Specifically, coengagement graphs represent popular users, hereafter referred to as influencers, as nodes in a network, and edges as representing a relationship where a large number of other users have retweeted both influencers repeatedly over a given time period. Edges between two nodes are weighted by the number of users who have retweeted both influencers at least n times. For this study, in order to remove users whose popularity had declined since the 2020 United States election, we constructed a new coengagement graph for the time period studied. In this graph, influencers were included as nodes if at least 500 other users retweeted both them and another influencer over the studied time period. This is process is meant not only to identify *popular* influencers, but popular users who share *community* with other influencers. This new graph reduces the studied influencer set to the 1,848 influencers referenced above.

To identify ideological communities among influencers, we run the Louvain clustering algorithm on the resulting coengagement graph, as implemented in the software package Gephi (Blondel et al., 2008; Bastian, Heymann, and Jacomy, 2009). We assign ideological labels by inspecting the usernames, user descriptions, and popular tweets of the most central users in each cluster. This process results in 988 mainstream progressive influencers, 757 pro-Trump influencers, and 61 radical socialist influencers. We provide further descriptions of these ideological clusters and justifications for our choice of labels in the Findings.

6.6.2 Labelling Quote Tweets

The primary indicator for negativity in this study is the quote tweet. The quote tweet is an affordance of Twitter that allows users to add commentary to another user's post, and rebroadcast both the original user's post and their commentary to their own audience. While seemingly a tool of convenience, it quickly became known (at least in the United States) as a tool for explicitly *denigrating* other users' posts (Schwedel, 2017; McNear, 2018). To study quote tweets, we extract 1.4M

unique quote tweets between our tracked influencers, collectively resulting in 130M downstream retweets by other users.

To train a model to classify quote tweets as having either a negative or non-negative stance towards the author of the quoted content. A negative stance is defined as indicating either disagreement with the author's statement, or any other disparagement of the author outside of the content of the text. We do not distinguish between positive and neutral stances, as in other stance detection tasks, because our coders could easily separate between these categories due to the brevity and ambiguity of the quote tweet format. This coding scheme also crucially differs from most other stance detection tasks, as we detect stance orientation towards an author, rather than the subjects about which an author is talking about. For example, a user could quote another user's disparagement of a politician, expressing further disparagement of that politician. While they would both be expressive a negative stance towards the *subject* of the author's text, the quoting user has a non-negative, agreeing stance towards the text's author. We make this modification to better expose social relationships between Twitter users, rather than create a map of valences towards different political topics.

Our labelling team consisted of nine different coders, who in total made 5,879 label assignments across 3,686 tweets sampled from the full dataset timespan. We held periodic meetings to review disagreements in coding and clarify our codebook. We also coded in pieces of two months at a time across the two years of data, and used a topic modeling algorithm on the full datasets of tweets from these time periods to help us understand relevant political events during coding (Grootendorst, 2022). Specifically, the coding team would review news articles relating to the most prevalent topics in each time period and discuss the meaning of posts referencing these topics.

We specifically coded instances of tracked influencers quote tweeting other tracked influencers, and followed a varied sampling strategy. We used a mixture of random sampling from different time periods, and sampling the most-retweeted quote tweets between tracked influencers in our dataset. In an attempt to intentionally code ambiguous texts, we ran an early version of the model

on a portion of our coded tweets, and then coded tweets which were classified as "negative" between two influencers who were ideologically aligned. Intuitively, negative quote tweets *should* be less likely between aligned influencers, and thus these tweets should be more likely to be either mistakes or highly discriminative.

For model analysis, we only kept tweets that at least two coders had coded, leaving 2,059 full tweets. When two coders disagreed, a third coder acted as a tiebreaker. To evaluate the reliability of this coding scheme, we calculate a Krippendorff's alpha of 0.72 over tweets with multiple coders (0.71-0.73, 95% confidence interval over 1,000 bootstraps of 90% of the coded data, Hayes and Krippendorff, 2007).

6.6.3 Classifying Quote Tweets

We use our labelled data to finetune a pretrained language model for our negative/non-negative classification task. Specifically, we finetune the PoliBERTweet algorithm, which was trained from dataset of tweets relating to the 2020 United States presidential election (Kawintiranon and Singh, 2022). PoliBERTweet was itself finetuned from the more general BERTweet (Nguyen, Vu, and Tuan Nguyen, 2020), an language model trained a larger dataset of tweets on multiple topics, which itself uses the BERT language model architecture trained with the RoBERTa training method (Devlin et al., 2019; Nguyen, Vu, and Tuan Nguyen, 2020). We preprocessed our data to remove specific Twitter usernames and URLs, and then trained the PoliBERTweet algorithm for 10 epochs using 10% of our data as held-out data using the Huggingface machine learning package (Wolf et al., 2020). We achieved an F1 score of 0.80 on a randomly sampled out-of-training dataset with 50 tweets predicted to be negative and 50 predicted to be non-negative, with a higher accuracy on non-negative tweets than negative tweets.

6.6.4 Regression Analyses

We use regression analysis to predict the percent change in follower count for influencer i from timepoint t to timepoint $t+1$, given information about their account and their incoming engagement at timepoint t. Timepoints are measured in months, although separate models are run with week- and day-long timepoints to test the sensitivity of this interval choice.

We use a simple linear regression, and report standardized and unstandardized coefficients using the *lm.beta* package (Behrendt, 2023). Count variables are logged before being input into the model. We use the ratio of retweets to likes and replies to likes as independent variables representing the distribution of incoming engagement. Engagement is defined as the sum of likes, replies, and retweets. Negative and non-negative quote retweet counts are derived from the model classifications described above, and in-group and out-group classifications are derived from the network analysis described above.

6.7 Findings

6.7.1 Ideological cluster descriptions

Our clustering method finds similar ideological groupings to (Beers et al., 2023a), suggesting that such clusters have remained stable from 2020-2023 (Figure 6.1). This influencer community is characterized by two large clusters, which we term the pro-Trump and mainstream progressive clusters, and one cluster less than a tenth of the size of either which we call the socialist cluster.

The pro-Trump cluster is characterized by influencers, alternative news websites (Starbird, 2017), and far-right elected politicians who generally align themselves with former president Donald Trump. While this community's Twitter engagement during the 2020 election was dominated by Trump's own Twitter account, his suspension and subsequent refusal to return after being unsuspended makes him absent in this graph (Twitter, 2021; LeBlanc and Duffy, 2022). No single account has achieved a similarly dominant status in his absence. Some of the top 20 most central

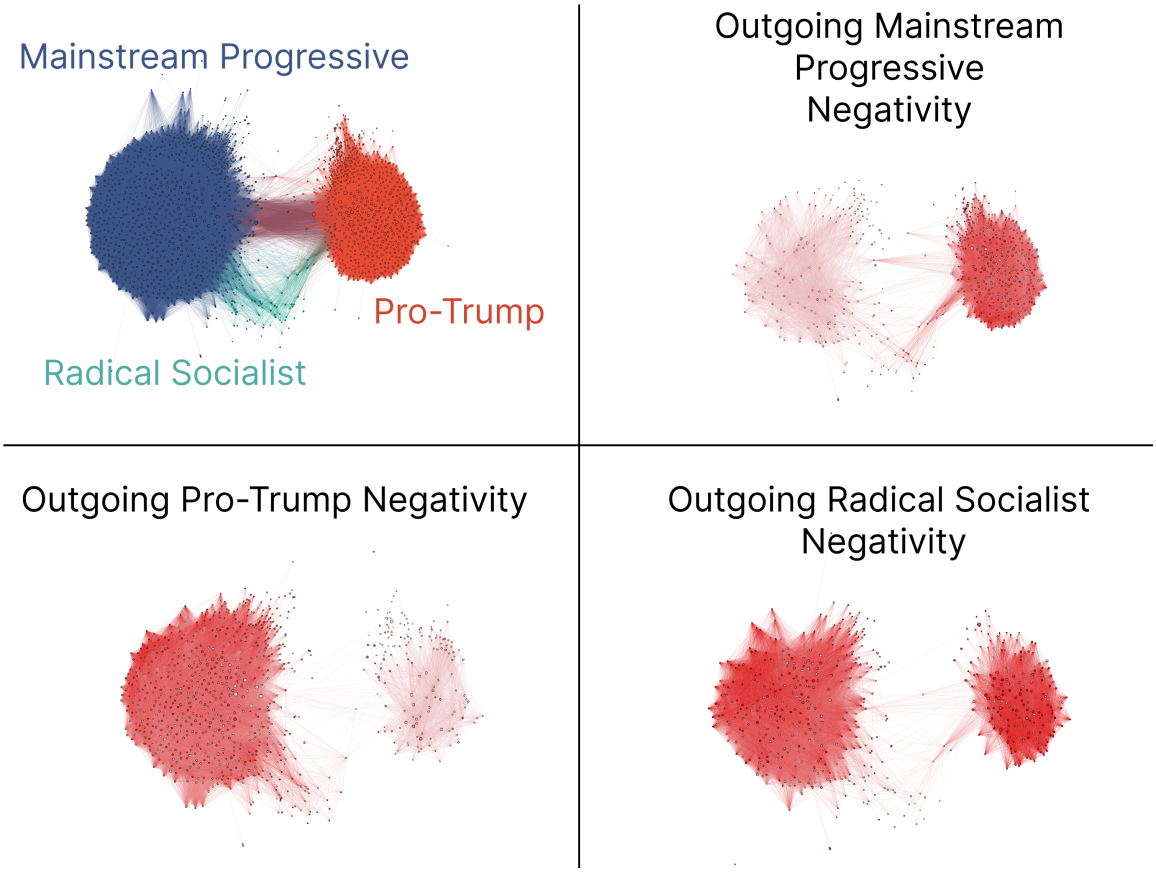

FIGURE 6.1: Panel A shows a visualization of a coengagement network, where each node is a Twitter user and each edge represents two users who share an audience of at least 500 users who retweeted them at least 15 times over the course of our dataset. Edges are weighted higher if more users retweet both nodes, and edges with higher weights are drawn more closely together. Nodes are colored by their cluster assignment via the Louvain clustering algorithm, and sized by their weighted degree. In panel B, three recolorings of the graph in Panel A are shown. Each represent the targets of negative quote tweets from the mainstream progressive, pro-Trump, and socialist clusters identified in Panel A. In these panels, nodes are sized by the number of times negative quote tweets of a node have been retweeted, and colored by the percent negativity of incoming quote tweets from this cluster, with deeper red nodes indicating higher overall negativity.

nodes in this cluster, according to weighted degree, include accounts for right-wing conspiracy theory influencer Jack Posobiec (*Jack Posobiec* 2022), Trump's son Donald Trump Jr. (Gray, 2024), far-right politician Lauren Boebert (Cameron, 2023), and Germany-based English-language disinformation news outlet Disclose.tv (Schumacher, 2022). While this cluster contains some accounts who are right-leaning but aligned (willingly or not) against Donald Trump, we will see that these accounts as are subject to consistent internal criticism for their perceived anti-Trump stances. Accordingly, we term this the "pro-Trump" rather than the, e.g., "right-leaning" or "Republican" (U.S. right-leaning party) cluster.

The mainstream progressive cluster is characterized by a similar array of left-leaning influencers, politicians, and media companies. For example, the top 20 most central nodes in this cluster include accounts for anti-Trump media organization the Lincoln Project (Williams, 2020), former U.S. Secretary of Labor and current multi-platform progressive influencer Robert Reich, and U.S. congressman Eric Swalwell. This cluster also includes "mainstream" media outlets, such as the New York Times and the Wall Street Journal, whose nominal commitments to objectivity sometimes preclude them from taking an explicitly progressive stance. However, similar to the pro-Trump cluster, these mainstream news accounts are frequently subjected to internal criticism when they create headlines counter to progressive stances. This cluster includes accounts for users who are both popular and frequently opposed to current president Joseph Biden or his policies, such as democratic socialist congresswoman Alexandria Ocasio-Cortez. Accordingly, we title this cluster the "mainstream progressive" cluster rather than the pro-Biden cluster, due to its range of progressive opinions and tolerance for disagreement with Biden's policies.

Finally, there is a much smaller cluster of what we term radical socialist accounts. These accounts include influencers and some media groups, but no elected politicians and few potential candidates. Accounts in this cluster generally take positions associated with socialist and radical U.S. politics, such as defunding or abolishing policing, making healthcare free and accessible, and opposing U.S. support of Israel. While these positions are sometimes found in the mainstream

	Quote Tweets	Quote Retweets	Negative Quote Tweets	Negative Quote Retweets	Negativity Multiplier
Overall	1.4M	130.3M	18.1%	15.0%	0.80x
In-Group	1.1M	99.7M	11.5%	7.0%	0.58x
Out-Group	0.3M	30.6M	39.2%	41.2%	1.09x

TABLE 6.1: Rates of quote tweet negativity, and quote tweet retweet engagement, as classified by our negativity model. The "Negativity Multiplier" column refers to the retweet performance of negative quote tweets versus non-negative quote tweets.

progressive cluster, they are intensified in the radical socialist cluster. Influencers in this cluster demonstrate a distrust for existing U.S. institutions, including elected politicians who identify themselves with socialism. They are notable for their opposition, in varying strengths, to the U.S. supporting Ukraine in the Russia-Ukraine war — a position more aligned with the anti-Trump cluster than the mainstream progressive. This cluster contains both news that adheres to traditional reporting norms, e.g. Jacobin Magazine (*Jacobin – Bias and Credibility* 2024), and those that engage in conspiracist journalism more typical of pro-Trump influencers, e.g. The Grayzone (*The Grayzone - Bias and Credibility* 2024). In these ways, and in the network graph, they are notable for their ability to support (and critique) both mainstream progressive and pro-Trump causes (Beers, Wilson, and Starbird, 2022).

6.7.2 RQ1: The use of negativity, within and between groups

We classified a total of 1.4M quote tweets made between our tracked influencers in these three clusters, collectively retweeted a total of 130M times (Table 6.1). Of those quote tweets, 18% were classified as negative, accounting for 15% of quote retweets, meaning that negative quote tweets garnered on average 0.80x as many retweets as non-negative quote tweets.

In general, groups send more quote tweets to in-group members, and get more retweet engagement when they do (99.7M in-group quote retweets, 30.6M out-group quote retweets). When group members did quote out-groups, those quote tweets were more likely to be negative (39%

negativity on out-groups vs 11% in-groups). These negative quote tweets of out-group members performed better than negative quote tweets of in-group members (1.09x average non-negative retweet totals for out-groups, but 0.58x for in-groups). Overall, negativity was more a frequent and effective strategy for engagement when targeted at out-group members than at in-group members, confirming hypotheses **H1a** and **H1b**.

6.7.3 RQ2: The targets of negativity

Unsurprisingly, influencers who received more non-negative attention (retweets, non-negative quote tweets) also received more negative attention, in absolute terms. By absolute retweet engagement, the most disliked out-group and in-group accounts tended to be those who were the most popular in general: president Joe Biden, the New York Times account, and far-right mega influencer Jack Posobiec are in the 25 most-disliked accounts by absolute retweet count. In general, the logged number of non-negative tweets has a 0.67 Pearson's R-squared correlation with the number of logged negative tweets they receive, suggesting that negativity is to a certain extent simply a side effect of receiving any attention at all. Like the distribution of retweet attention generally, the distribution of negative retweet attention is highly unequal, with the top 1% most-disliked influencers responsible for 25% of all attention in our dataset (Gini coefficient: 0.80).

Many influencers, however, are disproportionately targeted for negative attention relative to their overall attention, both within groups and between groups. To remove influencers who are targeted for quote tweets very infrequently, we focus here on influencers ranked in the top 20% for incoming quote tweets. To understand which accounts are more disliked than we would expect given their non-negative retweet and quote engagement, we generate a ranking of the ratio of negative engagement to non-negative engagement for each of these influencers. We find that the ratio of negativity influencers in this group is positively skewed, with most influencers receiving relatively low (<10%) proportions of negative quote engagement, and 10% of influencers receiving very high negative quote engagement (>50%). We find this applies for both negativity heading

towards in-groups and out-groups, validating hypotheses **H2a** and **H2b**.

Consistently, the most negatively-targeted accounts outside of a group are different from the most positively-targeted accounts within that group. Accounts which garner disproportionate negativity tend to have high public profiles off of Twitter, and particularly are disproportionately current/former elected officials or candidates. Table 6.2 shows that 9/10 of the most-disliked accounts by mainstream progressives are pro-Trump politicians, while 5/10 of the most-disliked accounts by pro-Trump influencers are. Only 3 of the most disliked accounts by radical socialist accounts are, perhaps due to their lessened focus on electoral politics.

Conversely, accounts which enjoy the most positive reputation among out-groups tend to be news accounts and non-self-published journalists. All ten of the least-disliked accounts by pro-Trump influencers are in this category, including news organizations like CSPAN and ABC. Similar trends hold for the mainstream progressive influencers (7/10 most-liked accounts) and radical socialist accounts (10/10). Mainstream progressive influencers notably had the lowest dislike ratios for conservative news accounts not directly associated with the U.S. pro-Trump movement, including British newspaper the Daily Mail, Israeli newspaper the Jerusalem Post, German disinformation outlet Disclose.tv, and seemingly non-partisan aggregator account Breaking911.

By contrast, the most disliked figures within in-groups are idiosyncratic, and tend to be popular users who violate a certain tenet of ideological orthodoxy within their cluster. For example, the top 10 most internally disliked pro-Trump accounts include, naturally, right-leaning accounts that have at one point expressed skepticism or criticism of former president Trump, tentative supporters of COVID-19 vaccination, outward supporters of Ukraine in the Ukraine-Russia war (such as Senator Lindsey Graham), and members of Republican congressional leadership (such as former congressman Kevin McCarthy), which during this time period was typically opposed to far-right, Trump-aligned politicians. Disliked mainstream progressive influencers, meanwhile, tended to oppose COVID-19 mitigation measures (such as journalist Nate Silver), or threaten president Biden's chance of re-election, in the case of third party candidate Andrew Yang. Radical

Highly Quoted Out-Group Accounts, by Quote Retweet Negativity					
Mainstream Progressive		Pro-Trump		Radical Socialist	
Account	Negativity Percentage	Account	Negativity Percentage	Account	Negativity Percentage
AndrewYang	80%	ewarren	82%	KamalaHarris	95%
tedcruz	80%	tedlieu	81%	RepSwalwell	93%
RonnyJacksonTX	79%	GeorgeTakei	80%	GOP	93%
Jim_Jordan	77%	JoyAnnReid	80%	laurenboebert	93%
MarshaBlackburn	74%	SenWarren	79%	DonaldJTrumpJr	92%
laurenboebert	74%	AOC	78%	robreiner	90%
benshapiro	74%	BernieSanders	78%	GeorgeTakei	89%
NikkiHaley	72%	davidhogg111	77%	ksorbs	87%
Mike_Pence	71%	SenSanders	76%	MattWalshBlog	87%
ksorbs	71%	cenkuygur	76%	benshapiro	87%

TABLE 6.2: The percentage of retweeted quote tweets classified by our model as negative between different ideological clusters. This measure sums the incoming retweets of all individual quote tweets, meaning highly retweeted quote tweets contribute more to the percentage. Quoting clusters are on the horizontal labels, while quoted clusters are on the vertical.

socialist influencers also tended to target fellow influencers who supported Ukraine in the Russia-Ukraine war, such as third party candidate Marianne Williamson.

We can also consider influencers who receive disproportionately little quote retweet attention for the non-quote retweet attention they receive. Figure 6.2 displays accounts with disproportionately high and low rates of negative attention via quote tweets and retweets combined. Accounts with high negative attention compared to their overall attention are once again likely to be elected politicians, like president Joe Biden or far-right senator Jim Jordan. However, this analysis reveals a second group of influencers who receive extremely high retweet engagement, but little to no negative quote tweet engagement. These influencers tend to be "platform-first," i.e. have little public notoriety outside of Twitter, but are incredibly active users who frequently attract high engagement on their posts. These include far-right conspiracy influencer Jack Posobiec, pseudonymous far-right influencer "catturd2" (Li and Kaplan, 2021), and progressive influencers Majid M. Padellan and Jon Cooper.

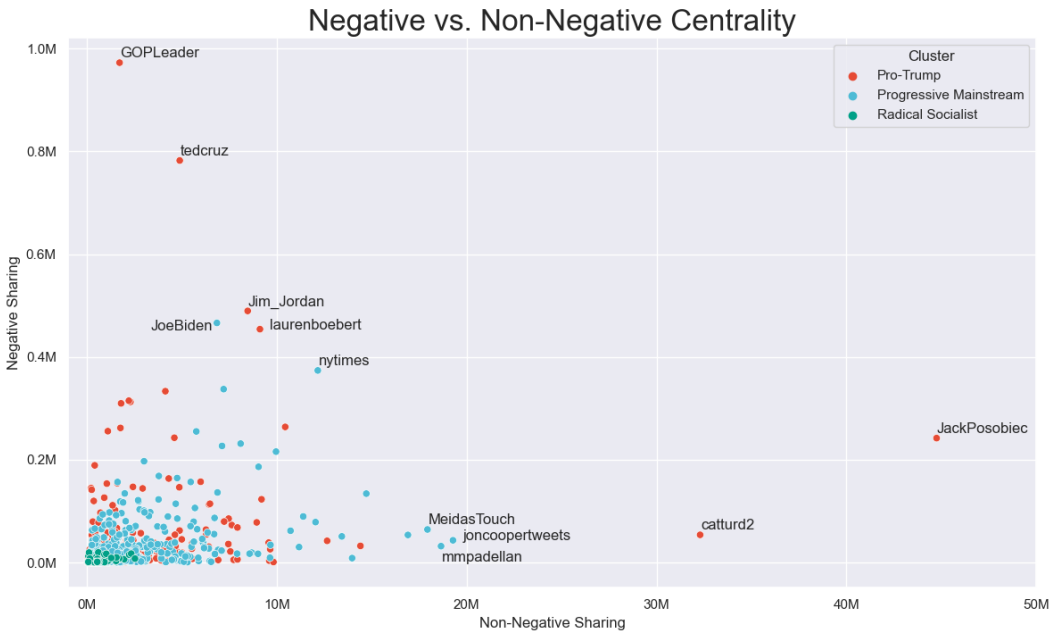

FIGURE 6.2: A scatter plot comparing negative with non-negative sharing of the influencers in our dataset. Negative sharing consists of retweets of negative quote tweets targeting a user, while non-negative sharing consists of retweets of a user's tweets and non-negative quote tweets targeting them. Only labels for influencers with 17M+ non-negative shares 0.35M+ negative shares are displayed to reduce visual clutter. All five of the left-most labeled accounts are elected politicians, while all five of the right-most are full-time influencers.

6.7.4 RQ3: The benefits and drawbacks of negativity

We now report findings from a regression analysis which attempts to predict influencer follower counts based on the reception of negative quote tweets, from in-groups and out-groups. Our first linear model regressed the percent change in influencer follower counts over month-to-month time periods against a variety of influencer-related predictors, and achieves an adjusted R-squared of 0.23 over 37,728 observations (Table 6.3).

The dominant effects in our model predicting an influencers' follower counts in the following period is the previous period's follower count and total engagements. Influencers with more engagement tended to grow their followers, but accounts with more followers need more engagement to maintain the same rate of growth. Holding the amount of incoming engagement constant, it was better in our data to have received that engagement on fewer tweets i.e. with higher average engagement per tweet. In general, making posts that garnered more retweets and replies relative to likes indicated stronger follower growth, although these effects are smaller than the main effects. Pro-Trump influencers in general saw stronger follower growth over time than other clusters, perhaps reflecting different following norms or different growth in audiences sizes in this cluster.

Regarding quote tweets, effect sizes are small. There are slightly benefits to being quoted non-negatively, but negative or zero effects of receiving negative quote tweets. We achieve nearly identical results in terms of coefficient sign and magnitude when training a model with day-to-day and week-to-week time periods instead of month-to-month time periods. In the day-long and week-long models, the effects of in-group negative quote tweet attention is slightly positive instead of negative, but regardless a very weak predictor. We reject both hypothesis **H3a** given that negative out-group targeting did not increase follower counts, and confirm weakly **H3b** that negative in-group targeting reduces follower counts. In general, however, we find that follower increases are much more dependent on one's *own* tweeting than on being quoted.

Variable	Coefficients	Standard Error	Standardized Coefficients
Intercept	3.94	0.2	NA
Engagements$_t$* **	0.92	0.01	0.5
Followers$_t$* **	-1	0.01	-0.46
Original Tweets$_t$* **	-0.32	0.01	-0.14
Retweets:Likes Ratio$_t$* **	1.6	0.23	0.04
Replies:Likes Ratio$_t$* **	2.75	0.34	0.04
Cluster: Radical Socialist	0.02	0.09	0
Cluster: Pro-Trump* **	1.27	0.03	0.18
Non-Negative Out-Group Quote Retweets$_t$* **	0.08	0.01	0.06
Non-Negative In-Group Quote Retweets$_t$* **	0.03	0.01	0.03
Negative Out-Group Quote Retweets$_t$* **	-0.04	0.01	-0.03
**Negative In-Group Quote Retweets$_t$* **	-0.02	0.01	-0.02

TABLE 6.3: Results of a regression model which predicts influencer follower counts at month $t + 1$ given dependent variables at month t, as well as cluster-level categorical variables. We take the log of all count variables before inputting into the model. Significant effects are bolded, with * indicating significance at .05 level, and *** indicating significance at .001 level.

6.7.5 RQ4: Ideological differences in the use of negativity

Different ideological clusters used quote tweets and particularly negative quote tweets differently (Table 6.4). In general, almost all quote tweets were aimed at the pro-Trump and mainstream progressive clusters, and not at the much smaller and lower-engagement radical socialist cluster. While groups had a tendency to quote tweet themselves more than other groups, the radical socialist group still only received 10% of its quote tweet engagement from in-group citations, compared to over 70% in the other two clusters. The radical socialist subgroup also did not split its quote tweeting equally between the other two groups, but rather quote tweeted mainstream progressives nearly twice as often as pro-Trump groups.

Pro-Trump and mainstream progressive influencers received roughly similar amounts of their total quote tweet engagement from negative quote tweeting, while radical socialist influencers received a slightly larger proportion from negativity (23.6% negative engagement for radical socialists vs. 16.6% for pro-Trump influencers vs. 13.9% for mainstream progressives). This is mostly due to their much higher targeting of out-groups, which inspire more negativity, than their own in-group. Pro-Trump and mainstream progressive influencers both hold more negativity for each other than radical socialists, and the highest negativity rate is found heading from mainstream progressives to Pro-Trump accounts (58.9% negative engagement towards pro-Trump, 32.6% vice versa). Thus, we can affirm hypotheses **H4a** and **H4b**: different groups use negativity to different extents, and target negativity differentially towards specific other groups.

The sign and relative magnitude of coefficients in the previous linear regression do not change if only run with one cluster at a time. Accordingly, we reject **H4c,** and conclude that the (lack of) benefits of being targeted for negativity are consistent across ideological clusters in this dataset.

Quote Retweet Distribution	Mainstream Progressive	Pro-Trump	Radical Socialist
Mainstream Progressive	82.8%	26.1%	59.2%
Pro-Trump	16.7%	73.6%	29.9%
Radical Socialist	0.5%	0.4%	10.9%

TABLE 6.4: The distribution of retweeted quote tweets between different ideological clusters. This measure sums the incoming retweets of all individual quote tweets, meaning highly retweeted quote tweets contribute more to the measure. Quoting clusters are on the horizontal labels, while quoted clusters are on the vertical.

Quote Retweet Negativity	Mainstream Progressive	Pro-Trump	Radical Socialist
Mainstream Progressive	7.0%	58.9%	12.7%
Pro-Trump	32.6%	6.8%	27.8%
Radical Socialist	24.0%	39.3%	18.3%

TABLE 6.5: Relative negativity of quote retweet engagement between different ideological clusters. Quoting clusters are on the vertical axis, while quoted clusters are on the horizontal axis. Values represent the percent of retweeted quotes which are negative when targeted from the quoting cluster to the quoted cluster.

6.8 Discussion

In this paper, we developed a model which could classify targets of quote tweets as "negative" — incorrect or disreputable — and used this model to conduct a large-scale analysis of influencer-influencer interactions in U.S. political discourse from 2021-2023 on Twitter/X. We show, via a network method which reveals political-ideological communities, that the use of quote tweet negativity follows expected patterns derived from social identity theory: namely, that in-groups tend to use negativity more often against out-groups than against their own members. However, the extent to which different ideological groups use negativity differs, with radical socialist influencers displaying particularly different usage patterns. Our qualitative analysis of highly-disliked accounts, especially accounts disliked by their own communities, suggests that negativity can be used to enforce norms within an online political community. Our findings on the disproportionately-disliked figures in out-groups suggests that out-groups are more likely to target users with high indicators of visibility, such as politicians, for negativity, and less likely to target news organizations and journalists. Our regression on the predictors of follower growth in

our dataset suggest that receiving negative quote tweet attention from in-groups or out-groups is not effective for follower growth, although *non*-negative attention from quote tweets is weakly effective.

Within the field of psychology, the social identity framework has suggested the negativity should be more frequent and endorsed when directed at out-group members, and conversely less effective when directed at (conforming) in-group members. Our results confirm these findings, and are in line with similar findings on quote tweets in political communities reported in Wojcieszak et al. (2022), and "hate-sharing" in far-right communities reported in Weigel and Gitomer (2024). Some popular media discourse has suggested a corollary to this theory: an influencer who attracts negative attention from out-groups will either gain influence within their own in-group, or possibly convert in-group members (McNear, 2018; Donovan and boyd, 2021). Our regression analysis suggests that, as a general practice among our tracked influencers, being the target of negativity is not effective for increasing follower counts, in-group or out-group. Non-negative attention, by contrast, is, particularly via direct retweets of one's original posts but also weakly via retweets of non-negative quote tweets of one's posts. Within our dataset, it is better for the aspiring influencer to be liked than disliked.

Our qualitative analysis of *how* negativity is used suggests that rather than as a tool for building one's following, negativity is used as a tool to enforce community norms. This is broadly consistent with previous work in communications (Gray, 2021), celebrity studies (Abidin, 2019; Christin and Lewis, 2021; Marwick, 2021; Lewis and Christin, 2022; Harris, Foxman, and Partin, 2023), journalism (Donovan and boyd, 2021; Miller, 2023), and cultural studies (Nakamura, 2015; Clark, 2020), which understand the use of negativity of and by influencers as a public performance which enforces or reaffirms group norms *and* group identities while punishing outsiders. Our qualitative findings of how influencers use negativity to chastise and stigmatize other influencers based on violations of in-group dogma reinforce these findings. Given the lesser engagement found on negative quote tweets targeting in-group members, it is plausible that negativity is used *specifically* to

enforce norms even the expense of engagement and follower growth.

Out study of multiple ideological communities of influencers displays key asymmetries in the way that quoting and negativity is applied between different groups, suggesting different group norms. We find two salient points worth noting. The first is that the over-representation of well-resourced journalistic organizations in the progressive mainstream cluster could plausibly create a dependence on this cluster. These far-right and radical socialist influencer communities can be considered examples of new, influencer-based alternative media systems. Audiences for alternative media historically display distrust in mainstream media (Cushion, McDowell-Naylor, and Thomas, 2021), but still use it selectively (Thorbjørnsrud and Figenschou, 2022; Andersen, Shehata, and Andersson, 2023). In a sense these alternative media systems are dependent on the mainstream media ecosystem, and limited in the amount of negativity they can treat it with. No such limitation exists in the opposite direction, letting adherents of mainstream media treat alternative media with greater negativity.

The second salient point regards the particularly outward orientation of the radical socialist cluster's retweet attention. This distinguishes it from the pro-Trump cluster, which while "alternative" to the mainstream cluster, is of equal size and supports and pays more attention to its own members than out-groups. It accordingly may be valuable to characterize the radical socialist cluster as a counterpublic, whose primary function is to deliberate and then advocate for change in discussions within the dominant public (Fraser, 1990; Squires, 2002). It follows that, unsurprisingly, negativity has different meanings depending on whether it originates from dominant/defensive publics or counterpublics: holding to account on the one hand, and enforcing dominance on the other (Jackson and Kreiss, 2023). We encourage additional research on how online counterpublics use targeted negativity to shape dominant public discourse and achieve policy change.

Prior research on morally motivated networked harassment, and harassment in general, has stressed the identity characteristics of harassment online. People attack each other not simply for differences in beliefs, but for differences in identities, like race, gender, sexuality, ability, and a variety

of other intersections and axes. Given this previous research, for example, it is suggestive that the six most-disliked accounts by our dataset's pro-Trump community, which is primarily a white nationalist, misogynist subculture, are women and people of color. We do not have the ability in this study to classify influencers by characteristics such as gender and race to understand how these factors may affect the disproportionate distribution of negativity — although we have ample evidence from prior research that such disproportionate harassment is occurring (Gray, Buyukozturk, and Hill, 2017; Lawson, 2018; Usher, Holcomb, and Littman, 2018; Harry, 2021; Thakur and Madrigal, 2022). However, it is likely that the unexplained differences in negativity targeting is in large part explained by this missing analysis of identity, and thus that an understanding of identity is key to understanding how negativity functions strategically in influencer communities. And even more, identity-based negativity by influencers may well serve a performative, discourse-shaping role by disciplining the public images of their targets into "controlling images," which limit the reach of marginalized communities in broader U.S. society (Steele, 2021).

6.8.1 Limitations

This paper has (at least) three clear limitations. The first is the limited accuracy of the classification algorithm. While effective enough to provide some insight into the dynamics of influencer negativity in a general sense, there are many instances in which our algorithm will have made errors. Given the algorithm's sole reliance on written language, it may be possible that certain topics or users' tweets are consistently judged incorrectly relative to our coders' intent. The potential for biased classification results could lead us to ignore or misinterpret certain data for qualitative analysis, and lead to misleading results in our quantitative analyses — particularly in linear regression. Future work could mitigate some of these risks by using regressions (and other tests) that explicitly model systematic bias introduced by classification errors (TeBlunthuis, Hase, and Chan, 2024).

The second is that the influencer set is limited due to its creation in the time and context of the

2020 United States election. We can identify members of this community likely to be missing by identifying who members of our dataset most often quote *outside* of our dataset. This analysis reveals some figures who have come to prominence in U.S. politics since 2020, including right-wing X owner Elon Musk (Warzel, 2022), the official presidential account for Joe Biden as opposed to his personal account, popular anti-transgender activist Chaya Raichik (Lorenz, 2022), and various popular right- and left-leaning full-time influencers. To the extent that these missing figures embody developing discourses not fully described by other users (such as anti-transgender politics), our analysis is incomplete.

The third is that we treat influencers as a single type, when they may in fact be composed of different types for whom negativity has different rewards. It is plausible, for example, that politicians, journalists, and media organizations follow different strategies from e.g. full-time influencers in order to gain followers, and may have more or less to gain by giving and receiving negativity. This is also a study of one particular community, where, for example, factors like political civility norms may disincentivize the use or reputational benefits of negative speech (Jamieson et al., 2017). Influencers in other communities for whom harassment is a stronger norm, such as the well-studied case of misogynist communities stemming from U.S. video game cultures (Gray, Buyukozturk, and Hill, 2017; Jamieson et al., 2017; Mortensen, 2018), surely face different incentives with respect to negativity. Future analyses should understand how both the type of community and influencer can affect the strategic benefits of using negativity online.

6.9 Acknowledgements

Funding for this work has come from the University of Washington's Center for an Informed Public, the John S. and James L. Knight Foundation (G-2019-58788), the Election Trust Initiative, and the National Science Foundation (grant #1749815 and grant #2120496).

Chapter 7

Study 4: Dueling Consensuses

Contributing Authors: Anna Beers[a], Sarah Nguyễn[b], Jevin West[b], Kate Starbird[a], Emma S. Spiro[b]

[a]University of Washington, Human Centered Design and Engineering

[b]University of Washington, Information School

7.1 Preamble

This chapter was originally published in Beers et al. (2023b), and the text and figures outside of this preamble section are unmodified. This preamble serves as additional contextualization for this study within the context of this dissertation, and personal reflection on my part for where this work has since travelled. I use single-spacing in this preamble to offset it from the main text.

This work took three and a half years to complete. It began as a rejected paper submitted to CSCW on boundary-work, and aspects of its analysis were incorporated into two other conference presentations (Beers et al., 2021a; Beers et al., 2021b). Aspects of it were also incorporated into unpublished work, including a still-unpublished dataset of 5,000 hand-coded tweets related to this dataset and an extensive side-project related to computational image analysis of COVID-19 memes. This is all to say: the research question in this paper derived from a particularly long series of failures, as I worked my way closer to the core of what this paper was really about.

That core ended up being about how science was used by non-experts online, and ultimately how existing methods to evaluate controversy in science did not accurately capture how controversy travelled through digital spaces. Existing literature on scientific controversy and misinformation often focused on scientists as the ultimate creators of information, and problems of science communication ultimately lied either with them or with the miscommunication of their work. This paper shifts the act of cultivating knowledge downstream to influencers and other media figures,

arguing that the act of communication ultimately is that act which creates consensus. It accordingly shifts the methods to what is precisely happening during science communication.

This study most directly addresses the third research question of this dissertation, how the activities of influencers shapes political discourse, although even then it is a somewhat indirect analysis. One has to begin by stating that seemingly medico-scientific discourse around COVID-19 *is* a political discourse, an intervention which is not explicitly made in this paper due to its short format. But I argue that the processes which apply to this discourse apply to all discourses that are animated by *evidence*, and that influencers play a key role in presenting and organizing evidence for their audiences while defending their audiences against threatening counter-evidence. They cut the information world in two, applying a positive/negative binary over hundreds of difficult-to-interpret scientific articles, and suggesting offline actions as a result of that binarization.

This is the only study which momentarily steps out of Twitter data, and arguably is a "multi-platform" study. The alternate data sources considered are mostly personal websites and blogs, news websites, and YouTube and other video-hosting services. The collection of this data was intensely manual, and yet very rewarding, providing me with a type of data that no one else had access to, and which could generate unique results. The writing of this paper is partly why I am not worried about wading into the post-API age. Manual annotation can quickly generate very interesting results, if you know wear to look, and I credit the process of writing this dissertation and exploring "big" datasets with giving me plenty of places to look.

I am currently following up this work with a longitudinal paper using a similar study design. I had originally designed this study with a longitudinal format, attempting to study how alternate consensuses emerge from nowhere. The mask case seemed like a particularly fruitful example, given how few non-experts cared about masking before the COVID-19 pandemic. However, the data ended up being too sparse to draw any conclusion in particular. The new study takes the form of a ten-year development of an alternate consensus around transgender healthcare, specifically ones that argues against the provision of that healthcare. This new study also has the advantage of addressing a medical discourse which is much more obviously politicized, and for which the "mainstream" scientific perspective is still deeply flawed. One of my struggles in this paper was figuring out how to castigate conspiracy science while still not launching a defense of (often harmful) mainstream science — something which will be easier to do when the bad transgender medicine of today is the good transgender medicine of yesterday.

7.2 Abstract

The COVID-19 pandemic provides a unique opportunity to study science communication, and in particular, the transmission of consensus. In this study, we show how "science communicators" — writ large to include both mainstream science journalists and practiced conspiracy theorists —

transform scientific evidence into two dueling consensuses, using the effectiveness of masks as a case study. We do this by compiling one of the largest, hand-coded citation datasets of cross-medium science communication, derived from 5M Twitter posts of people discussing masks. We find that science communicators selectively uplift certain published works while denigrating others to create bodies of evidence that support *and* oppose masks, respectively. Anti-mask communicators in particular often employ selective and deceptive quotation of scientific work, and criticize opposing science more than pro-mask communicators. Our findings have implications for scientists, science communicators, and scientific publishers, whose systems of sharing (and correcting) knowledge are highly vulnerable to what we term adversarial science communication.

7.3 Introduction

Public consensus often diverges from scientific consensus (Lewandowsky, Gignac, and Vaughan, 2013; Scheufele and Krause, 2019). Something is lost in translation from the scientist to the public, as "science communicators" — writ large to include both mainstream science journalists and practiced conspiracy theorists — reshape and transform scientific evidence in ways that may deviate from its authors' intents (Oreskes and Conway, 2012). For example, in 2015, a paper suggested that cloth masks increase the risk of respiratory illness in hospital contexts (MacIntyre et al., 2015). It became one of the most shared papers about masks on social media during the COVID-19 pandemic, was quoted in one of the most popular television news programs in the United States, and was cited in legal challenges to mask mandates (Hitlin, 2016; *Tucker Carlson Tonight* 2021). However, the paper's first author wrote more than 10 papers over a two year period during COVID-19 affirming the benefits of masks and speaking out against the selective and misleading understanding of the 2015 paper, before and after its popularity on anti-mask media (MacIntyre and Chughtai, 2020a; Chughtai, Seale, and Macintyre, 2020; MacIntyre et al., 2020; MacIntyre and Chughtai, 2020b; Wang et al., 2020; Chughtai et al., 2020; Raina MacIntyre, Costantino, and Chanmugam, 2021; MacIntyre et al., 2021; MacIntyre and Binkin, 2021; MacIntyre et al., 2022).

How does this re-imagining of scientific consensus outside of the hands of scientists themselves happen in today's media-rich environment? What are the implications, and what can be done about it?

In this paper, we hand curate one of the largest corpora for investigating these questions. Leveraging the high demand for research findings about masks during the early stages of the COVID-19 pandemic, we show via citation analysis how science communicators and practiced conspiracy theorists alike manufacture the appearance of consensus, using selective and critical citation to create dueling, opposite perceptions of scientific knowledge. We find that scientific publishers and authors are unprepared for how their work is being misused by "science communicators" to advance scientific misunderstanding.

7.3.1 Science communication during a pandemic

Rarely have so many people been so personally invested in emerging scientific discoveries, and never before has the process of science communication been so thoroughly documented on social media as it has during the COVID-19 pandemic. These circumstances have shown vividly how varied the sources of scientific knowledge have become: not just scientists, but journalists (Brüggemann and Engesser, 2017), government officials (Lee and VanDyke, 2015), celebrities (Alatas et al., 2019), viral videos (Allgaier, 2019), and social media users themselves (Lewandowsky et al., 2019; Lee et al., 2021). We present and analyze here two of the largest hand-coded datasets of scientific citation to understand the development of a pandemic controversy around masks in the United States (US). The first dataset shows how scientists built and shared amongst themselves a varied evidence base on the effectiveness of masks for COVID-19. The second dataset shows how science communicators, credentialed and not, transformed this evidence base into two different, opposing perceived consensuses: one in favor of masks, and one against. In documenting this process of consensus transformation, we have three primary findings:

- While scientists largely agreed that masks were effective at reducing the spread of infectious

disease and tended to cite works with similar *methodologies*, science communicators sharply disagreed and tended to cite work with similar *stances*.

- While scientists rarely engage in negative citation, science communicators and particularly anti-mask science communicators more frequently disparage and criticize scientific work.
- Science communicators cite much of the same literature as scientists. However, consensuses diverge via misleading contextualizations that often contradict the scientists publicly stated views.

While popular rhetoric has held that those who disbelieve mainstream science are "anti-science" (Tyson, 2021; Philipp-Muller, Lee, and Petty, 2022), our dataset describes a communicative world where both pro- and anti-mask social media users are attuned to scientific norms and invest heavily in citing the primary literature (Koltai and Fleischmann, 2017). Indeed, in what is akin to context collapse (Marwick and Boyd, 2011), anti-mask science communicators gladly repurpose the knowledge infrastructure of mainstream science to spread misinformation in the language of science itself (Pasquetto et al., 2022). These science communicators' frequent use of negative citation (Catalini, Lacetera, and Oettl, 2015) can be seen as inoculating their audiences against the views of mainstream science – the mirror of a strategy of "prebunking" and "inoculation" advocated by science communication researchers (Linden, Roozenbeek, and Compton, 2020; Lewandowsky and Linden, 2021; Basol et al., 2021). Anti-mask science communicators' deceptive but likely effective citational practices reveal how scientific authors and publishers are currently woefully unprepared for the ways that their articles are misrepresented to the public (Boutron and Ravaud, 2018).

Conceptually, this work bridges consensus research in science communication, which focuses on the public communication of science (Kahan, Jenkins-Smith, and Braman, 2011; Lewandowsky, Gignac, and Vaughan, 2013; Chinn, Lane, and Hart, 2018; Stekelenburg et al., 2022); science and technology studies, which focuses on how consensus is constructed and then understood by both scientists and the public (Marres, 2015; Miller, 2019; Oreskes, 2021); and the science of

science, which focuses on consensus formation through citational networks (Edge, 1979; Bruggeman, Traag, and Uitermark, 2012; Mingers and Leydesdorff, 2015; Fortunato et al., 2018). We focus on the role of formal (government health bulletins) and informal (anonymous viral Twitter posts) science communication. The work is inspired by the concept of knowledge infrastructures — the "robust internetworks of people, artifacts, and institutions which generate, share, and maintain specific knowledge about the human and natural worlds" (Edwards et al., 2013). By centering both formal and informal knowledge creators as objects of analysis, we create new opportunities for quantitative and qualitative inspections of knowledge infrastructures and the way they shape public health outcomes.

7.4 Results

7.4.1 The creation of evidence bases by scientists

In these findings, we make a distinction between scientists, who publish their work for consideration by the scientific community, and science communicators, who publish their work for consideration by the lay-public and may or may not include formally trained scientists. We assume that the goal of science communicators in crises is to translate the (sometimes developing) consensus of scientists into public-facing information about probable risks and recommended responses (Fischhoff and Scheufele, 2013). In practice, we find that popular science communicators in the US are sharply divided on the benefits of masks for preventing COVID-19. Thus, our first task is to understand whether these science communicators are faithfully transmitting a true dissensus among scientists themselves, or are generating it from whole cloth.

To assess the existing evidence base on mask-wearing at the start of the pandemic, we use the Web of Science (Analytics, 2017) to identify papers about masks that were highly-cited in the first year of the COVID-19 pandemic. We annotate all citations in these papers, and mark if they are positive/neutral citations, or negative citations. We define negative citations as citations that go

beyond discussing a work's limitations or conflicting results to actively recommend that readers at least partly disregard the cited work. We construct a signed (i.e., positive and negative citations specified through hand coding) co-citation graph of these highly-cited papers, in order to identify subsets of papers that tend to be cited together (Figure 1). We proceed under the assumption that papers frequently co-cited *positively* together, within a topical area, represent endorsements of shared evidence on the topic of masks in COVID-19 (Trujillo and Long, 2018).

We find three evidence bases that correspond to different methodologies and fields that highly-cited masks papers tend to reference. Specifically, these papers cite 1) studies grounded in particle physics that evaluate the material properties of masks, 2) studies grounded in epidemiology, which evaluate masks effectiveness on different metrics during the first year of COVID-19, and 3) studies grounded in public health, which evaluate the effectiveness of masks in different contexts, such as hospitals or households, and evaluate practical considerations such as the disinfection of masks. Each evidence base is composed of papers that mostly suggest that masks are effective for preventing respiratory illness, and especially COVID-19.

We find that in general scientists use negative citation sparingly, with only 1.3% of citations across 80 papers manually coded as negative (32 out of 2,426 citations). This percentage is notably lower than previous estimates, which have ranged from 2-15% (Moravcsik and Murugesan, 1975; Case and Higgins, 2000; Catalini, Lacetera, and Oettl, 2015). Affirming Bruggeman et al. (Bruggeman, Traag, and Uitermark, 2012), we find that negative ties, while rare, are not randomly distributed throughout the citation corpus. Rather, they are disproportionately directed at a few papers, including both papers with positive and negative results towards masks. The most negatively cited paper, with four negative citations, is MacIntyre et al.'s 2015 paper in *BMJ Open* that "cautions against the use of cloth masks" for preventing respiratory illness in hospitals, and suggests they may increase the risk of infection to the wearer (MacIntyre et al., 2015). This paper is particularly notable because its lead author has subsequently published no fewer than 10 publications over a

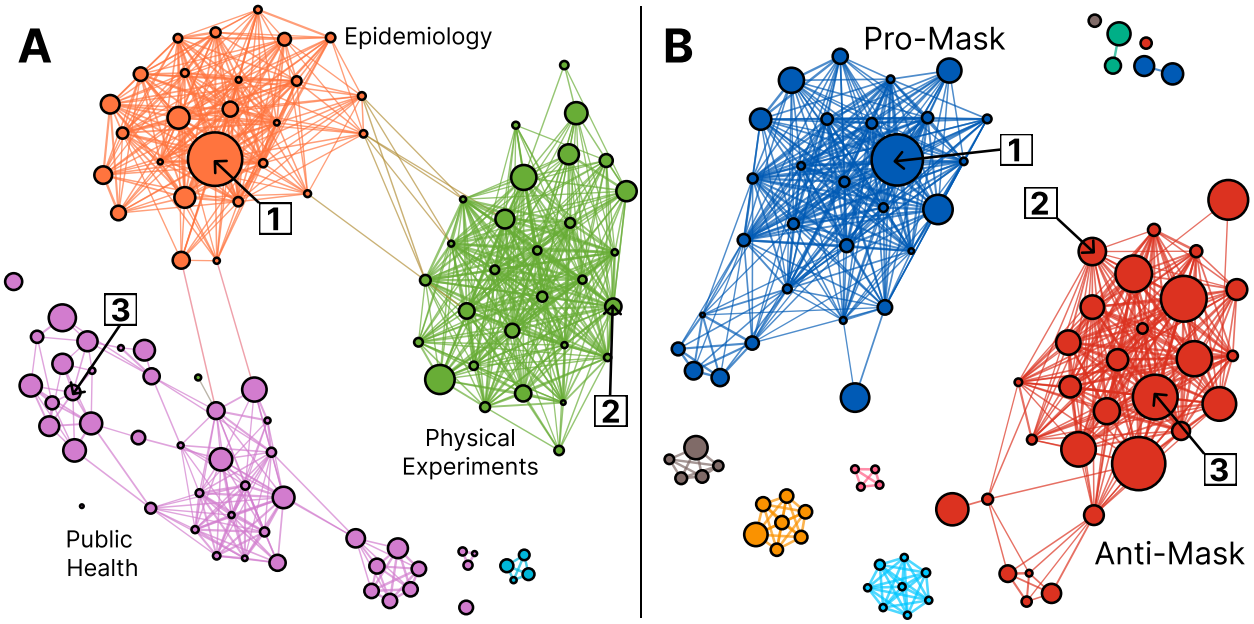

FIGURE 7.1: **A visualization of the cited literature of scientists and science communicators with respect to mask-wearing and COVID-19**. Papers that tend to be cited together by A) scientists and B) science communicators discussing masks and COVID-19. Scientists tend to cite groups of papers by topic or field. In A), the green cluster represents papers concerned with the physical properties of mask materials, the orange cluster represents epidemiological papers on masks effectiveness against COVID-19, and the purple cluster represents older public health papers on the effectiveness of masks in hospitals and households for COVID-19 and other respiratory illness. In B), meanwhile, science communicators cite largely the same papers as scientists, but segregate their citations by stance: supporting the use of masks in blue, and opposing them in red. Three papers (nodes) are labeled in both clusters: 1) (Leung et al., 2020), the most shared paper in both datasets, 2) (Rengasamy, Eimer, and Shaffer, 2010), a fluid physics paper from 2010 shared in the anti-mask perceived consensus, and (MacIntyre et al., 2015), a 2015 paper with a negative finding on masks whose lead author has since publicly supported mask-wearing for COVID-19. Larger nodes in the scientist network (A) are cited by more highly-cited papers on masks, while larger nodes in the science communication network (B) are cited by more frequently-shared Twitter URLs in a dataset of people arguing about masks.

two year period arguing for the use of masks of all types to lessen the spread of COVID-19 (MacIntyre and Chughtai, 2020a; Chughtai, Seale, and Macintyre, 2020; MacIntyre et al., 2020; MacIntyre and Chughtai, 2020b; Wang et al., 2020; Chughtai et al., 2020; Raina MacIntyre, Costantino, and Chanmugam, 2021; MacIntyre et al., 2021; MacIntyre and Binkin, 2021; MacIntyre et al., 2022). Among these publications is a letter titled "Masks in the community are an effective strategy," and a follow-up to the criticized 2015 study offering a possible explanation for their earlier, negative findings around cloth masks (MacIntyre and Chughtai, 2020a; MacIntyre et al., 2020). Chu et al.'s 2020 paper claiming that "face mask use could result in a large reduction in risk of [respiratory] infection" was also negatively cited three times primarily on methodological grounds, some of which are described by Jüni et al. in a letter response published in the same journal (Chu et al., 2020; Jüni et al., 2021). Overall, scientists negatively cite papers that have perceived methodological weaknesses, rather than papers that share a certain stance on the effectiveness of masks.

In this first year of the pandemic, the consensus on the effectiveness of masks was still developing, and new studies continue to be published since our analysis period. We show here not a proof of consensus that mask-wearing is an optimal choice for combating COVID-19, but rather the creation of a shared literature base from which scientists, within their own fields, collectively endorsed and used to move scientific knowledge and work forward. In the following section, we describe the different way that science communicators select literature bases, namely by their ability to support a pro- or anti-mask stance.

7.4.2 The creation of perceived consensuses by science communicators

Science communicators, by contrast, come to remarkably diverging conclusions about the effectiveness of masks within their own community. To surface the published works of science communicators, we analyze the URLs shared in posts by Twitter users discussing masks. These URLs sometimes link directly to the published scientific literature. But more often these URLs link to the

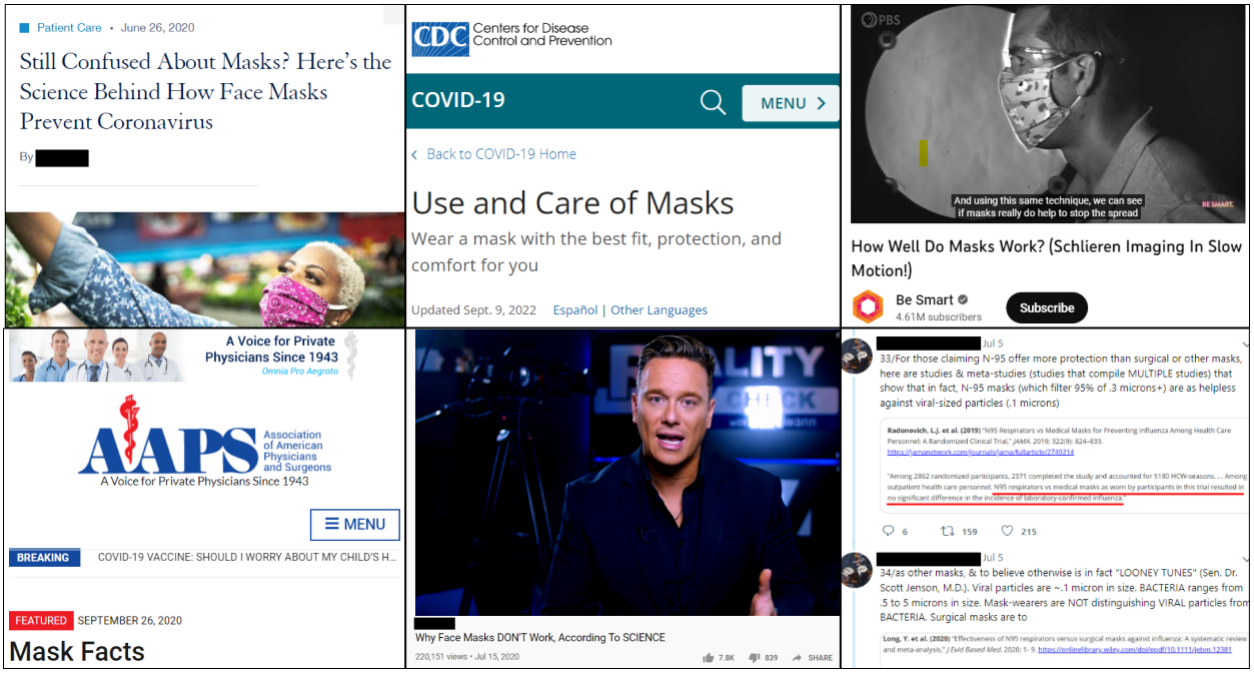

FIGURE 7.2: **Screenshots of six popular science communicators in our dataset.** The top row are examples of generally pro-mask science communicators, including from left to right an article published by the University of California San Francisco (UCSF), a help page published by the CDC, and a YouTube video published by the Public Broadcasting Service (PBS). The bottom row contains generally anti-mask science communicators, including from left-to-right a page published by the medical conspiracy group, the Association of American Physicians and Surgeons (AAPS), a YouTube video published by a conspiracy theorist formerly employed by Russian state media RT, and a viral threaded Twitter post by an anonymous user containing URLs and screenshots of scientific articles. We have appended black boxes to first image in the top row, and the second and third images in the bottom row, in order to protect the privacy of individual authors in these pieces

works of science communicators, broadly defined: science journalism articles, government- and hospital-sponsored websites, self-published personal blogs, viral social media posts, and videos hosted on a variety of platforms. Some examples of frequently-shared URLs which *support* mask-wearing include, with screenshots in the top row of Figure 2:

- "Still Confused About Masks? Here's the Science Behind How Face Masks Prevent Coronavirus," hosted by the University of California San Francisco (UCSF) Medical School (Bai, 2020).
- "Guidance for Wearing Masks" hosted by the United States Center for Disease Control (CDC) (Disease Control & Prevention, 2022).
- A video hosted on YouTube titled "How Well Do Masks Work? (Schlieren Imaging In Slow Motion!)" published by the United States Public Broadcasting Service (PBS) (Hanson, 2020).
- A long, threaded Twitter post which aggregates links to scientific studies purportedly supporting masks, created by a pseudonymous Twitter user. (nuanceORDEATH, 2020)

Conversely, examples of highly-shared URLs that oppose mask-wearing, again with screenshots in the bottom row of Figure 2, include:

- "Mask Facts," a website hosted by the Association of American Physicians and Surgeons (AAPS), a medical conspiracy group aligned with the US libertarian movement, which opposes the scientific consensus on vaccines, HIV/AIDS, climate change, and abortion among other topics (Singleton, 2020).
- "Masks Don't Work: A Review of Science Relevant to COVID-19 Social Policy," a self-published literature review authored by a scientist who also contested the scientific consensus on climate change and vaccines. This review has been re-hosted on many different websites, including the science aggregator ResearchGate (Rancourt, 2020).

- A video hosted on YouTube titled "Why Face Masks DON'T Work, According To SCIENCE" created by a journalist then employed by *RT*, a Russian state-owned news outlet known to publish false and misleading content. The author advocates for counter-consensus views on vaccines and a variety of conspiracy theories related to mass shootings in the United States, the Syrian Civil War, and other topics (Swann, 2020).
- "The Science of Masks," a website and companion video that links to over 60 studies and news articles about masks. It is hosted on the website of an alternative medicine podcast host that sells diet planning services, books about wellness, and speaking opportunities (Stevenson, 2021).
- A long, threaded Twitter post that aggregates links to scientific studies and other materials purportedly opposing masks, created by a pseudonymous Twitter user (jcho710, 2020).

The science communicators behind each of these websites, whether pro- or anti-mask, explicitly cite peer-reviewed scientific literature to explain their positions on the effectiveness of masks. Anti-mask science communicators cite the scientific literature in ways similar to pro-mask communicators, for example by using explicit "References" sections with formatted citations. They cite non-scientific sources, such as personal blogs or anecdotal news stories, only rarely. Most anti-mask communicators use sophisticated, scientific language to contextualize their findings, for example referring to "randomized controlled trials" and "statistical significance," and likely give the impression of a knowledgeable expert explaining the current state of scientific consensus.

To understand how science communicators reference science differently from scientists, we code their citations just as we did with the Web of Science data. We find that the proportion of negative citations is notably higher from science communicators than scientists themselves at 6% (89 out of 1,496). This proportion rises to 9% for websites opposing or reporting negative findings on masks, compared to 4% for websites with neutral or positive findings. Overall, anti-mask websites averaged 1.5 negative citations per website, while other websites average 0.7 citations.

Indeed, several anti-mask websites have sections specifically devoted to refuting research that supports masks. For example, one such section is titled "Popular Mask Effectiveness Studies That Ignore Real World Conditions," and lists the titles of pro-masks papers accompanied with short explanations of their supposed flaws in scientific language.

As with the Web of Science data, we use this dataset to create a signed co-citation network (Figure 1). We see several small groups of 2-7 papers, each likely the result of different science communicators using extremely similar citation lists, or in some cases directly copying and excerpting each others' pieces. The rest of the graph is dominated by two large groups of papers which, unlike citations in the scientific literature network, group together by *stance*. One group is pro-mask, and consists entirely of literature from mostly high-impact scientific journals whose results support the claim that masks are at least somewhat effective for preventing the spread of COVID-19 or other respiratory illnesses. A typical example of a paper in the pro-mask evidence base is Leffler et al. 2020 (Leffler et al., 2020), which states in its abstract that "societal norms and government policies supporting the wearing of masks by the public, as well as international travel controls, are independently associated with lower per-capita mortality from COVID-19."

The other group is anti-mask, and consists of the scientific literature and one university-affiliated website with positive, qualified, or negative findings about masks, almost all of which are construed to appear as negative findings in the context of their citations. A typical example in this cluster is Jacobs et al. 2009 (Jacobs et al., 2009), which argues in its abstract that "face mask use in health care workers has not been demonstrated to provide benefit in terms of cold symptoms or getting colds." Both groups of papers contain almost exclusively papers published in academic journals, and neither cluster display a clear preference for papers from the fields or journals identified in the scientific literature network.

In fact, both the pro- and anti-mask paper groups leveraged papers from the three different fields found in the scientific literature, extracting those which can be made to appear in agreement with

their chosen stance. Some papers are cited positively by both pro- and anti-mask science communicators: 15% of incoming citations to the pro-mask cluster come from websites with anti-mask orientations, and 23% of incoming citations to the anti-mask cluster come from websites with pro-mask orientations. However, rather than take these cross-stance citations as evidence of shared scientific knowledge between pro- and anti-mask perspectives, we find that papers are often contextualized differently between pro- and anti-mask websites (see Table 1 for excerpts). Accordingly, the separation between the evidence bases that readers perceive may be even wider than this data can illustrate. We discuss two such papers at length to illustrate two distinct tactics by which ostensibly pro-mask papers can be contextualized to appear to advise against masks: selective quotation and deceptive quotation.

We first examine Leung et al. 2020 (Leung et al., 2020), the paper which was most often cited by science communicators in our dataset. Inspecting how this paper was cited in-text, we find that pro- and anti-mask websites selectively quote entirely different parts of the paper. Pro-mask science communicators tend to quote or paraphrase the top-level finding from the abstract: "surgical face masks could prevent transmission of human coronaviruses and influenza viruses from symptomatic individuals." Anti-mask science communicators instead quote a sentence in the Discussion section, "our findings indicate that surgical masks can efficaciously reduce the emission of influenza virus particles into the environment in respiratory droplets, *but not in aerosols*" [emphasis added, for additional discussion of aerosols and COVID-19, see (Randall et al., 2021)]. In practice, Leung et al. exists as two papers with two, seemingly contradictory conclusions: one that claims that surgical masks can prevent transmission of COVID-19, and another which claims that they cannot, given that COVID-19 is an aerosolized disease. We observe that science communicators tend to selectively quote or cite the paper to support their stance on masks.

Next, we look at Long et. al 2020 (Long et al., 2020), a paper which is cited positively by both popular pro- and anti-mask websites. Long et al. concluded from their literature review on mask-wearing that "the use of N95 respirators compared with surgical masks is not associated

TABLE 7.1: A list of the top five most-shared citations (measured by number of shares of sources that cited them) in the Twitter dataset. We provide examples of positive and negative framings for each of the papers' results/conclusions.

Paper Citation	Pro-Mask Context	Anti-Mask Context
Leung et al. (2020) (Leung et al., 2020)	"It found that even loose-fitting surgical masks blocked almost all the contagious droplets the wearers breathed out and even also some infectious aerosols — tiny particles that can linger in the air." (Godoy, 2020)	"Surgical masks can efficaciously reduce the emission of influenza virus particles into the environment in respiratory droplets, but not in aerosols" (Singleton, 2020)
MacIntyre et al. (2015) (MacIntyre et al., 2015)	"One of the most frequently mentioned papers evaluating the benefits and harms of cloth masks have been by MacIntyre et al. Findings have been misinterpreted, and therefore justify detailed discussion here... Its implementation does not inform the effect of using cloth masks versus not using masks in a community setting for source control of SARS-CoV-2, which is of the same genus as seasonal coronavirus, which has been found to be effectively filtered by cloth masks in a source control setting." (Howard et al., 2020)	[*Transcribed from video*] "Now here's what's even more shocking about the study. Not only were the rates of infection higher in the cloth mask group than they were in the medical mask group. The rates of infection in the cloth mask group were significantly higher than the people in the control group, who are allowed to occasionally wear a mask or not even wear a mask at all." (Stevenson, 2021)
Long et al. (2020) (Long et al., 2020)	"Further, the evidence is mixed from randomized studies on types of masks and risk for influenza-like illness transmission to mask wearers; for example, a recent systematic review and meta-analysis comparing N-95 respirators versus surgical masks found a statistically insignificant decline in influenza risk with N-95 respirators." (Lyu and Wehby, 2020)	"There were no statistically significant differences in preventing laboratory-confirmed influenza, laboratory-confirmed respiratory viral infections, laboratory-confirmed respiratory infection, and influenza-like illness using N95 respirators and surgical masks." (Rancourt, 2020)
Radanovich et al. (2019) (Radonovich et al., 2019)	"Radonovich et al found in an outpatient setting that 'use of N95 respirators, compared with medical masks... resulted in no significant difference in the rates of laboratory-confirmed influenza'" (Howard et al., 2020)	[*Transcribed from video*] "In the 2019 study, there was no significant difference by wearing an N95 respirator mask." (Swann, 2020)
Smith et al. (2016) (Smith et al., 2016)	"If someone asks: What's the evidence for mask wearing? Here is a list of *SEVENTY* papers, including reviews/meta-analysis and individual studies, in reverse chronological order. Includes 31 from 2020 alone (!!). META = meta-analysis or systematic review." (nuanceORDEATH, 2020)	"This 2016 meta-analysis found that both randomized controlled trials and observational studies of N95 respirators and surgical masks used by healthcare workers did not show benefit against transmission of acute respiratory infections." (Huber, 2020)

with a lower risk of laboratory-confirmed influenza," and recommended that surgical masks may thus be a suitable substitute for N95 respirators. In three separate anti-mask websites, their results were quoted out of context as follows: "There were no statistically significant differences in preventing laboratory-confirmed influenza, laboratory-confirmed respiratory viral infections, laboratory-confirmed respiratory infection and influenza-like illness using N95 respirators and surgical masks." Taken alone, this quotation can make it appear that neither surgical masks nor N95 respirators had any effectiveness for preventing disease, a finding that directly contradicts the main result of Long et. al., which is about the effectiveness of N95 masks *compared to* surgical masks. In this case, quotation is not so much selective as deceptive, implying that surgical masks are ineffective through the removal of additional context, rather than that surgical masks are as effective as N95 masks – and both of which are associated with lower risk of illness.

We conclude our findings by noting that the anti-mask paper group is not a marginal phenomena in the realm of Twitter debates, and indeed the websites that cite them are shared more often than those that share the pro-mask paper group. To quantify the difference in how often Twitter users are exposed to these two paper groups via shared URLs, we sum the shares for each cited publication across all the science communicators citing them positively. According to this metric, nodes in the anti-mask cluster are shared positively more often than nodes in the pro-mask cluster (124K shares vs. 76K shares). In fact, by this metric, some nodes in the pro-mask cluster have more negative exposure than positive exposure, such as Leffler et al. 2020 (Leffler et al., 2020), which is not true for any node in the anti-mask cluster. We echo other researchers' caution of "digital bias" (Marres, 2015), namely that arguments offered in online forums are not likely to reflect public opinion writ large. Our chosen forum in this analysis, Twitter users arguing with each other in the wake of public health announcements by US governors, likely attracts a unique population subset of politically-engaged, internet-savvy commenters. In fact, anti-mask viewpoints may have been specifically elicited by the choice of a highly politicized discussion forum. However, to the extent that US governors' official announcements on Twitter are meant to be a public and visible forum

for discussing the pandemic, responses in the aggregate support the appearance of a misleading anti-mask consensus and oppose the mainstream pro-mask consensus.

7.5 Discussion

We describe how science communicators cite the collected works of scientists differently to create the appearance of two radically different scientific consensuses: one supporting mask use during COVID-19 and one opposing. We find that science communicators perform a transformation of the original evidence bases used by scientists, which are mostly organized by fields and methodologies, into separate evidence bases segregated by *stance* towards masks. Science communicators, who on one side are from mostly mainstream scientific and journalistic institutions and on the other are mostly a coalition of conspiracy theorists and creators on alternative media platforms, collect different scientific papers and through a process of aggregation (collection of positive citation) and denigration (negative citation) construct a perceived scientific consensus for their audiences. Some papers are referenced by both pro- and anti-mask science communicators, but rather than indicating shared scientific understanding, we see that these papers are either selectively or deceptively quoted to imply the opposite of their conclusions by anti-mask communicators. Overall, we portray the development of an anti-mask community of belief which, far from being "anti-science" as their detractors may claim (Philipp-Muller, Lee, and Petty, 2022), puts concerted (and deceptive) effort into the boundary-work of separating favored "good science" from disfavored "bad science" (Gieryn, 1983). We also contribute a method for mapping the transformation which occurs from the work published by scientists — so often the stopping point of scientometric analysis — into the consensus that the public perceives in online settings.

We find that negative citation is an important practice in the development of these two separate perceived consensuses on Twitter, and that anti-mask science communicators in our data use negative citation more frequently than other communicators. Science communication scholars have argued that prebunking, i.e. telling audiences that they will encounter misinformation,

may be an effective way to communicate science and inoculate them against false beliefs (Linden, Roozenbeek, and Compton, 2020; Lewandowsky and Linden, 2021; Basol et al., 2021). We find that, unfortunately, charismatic disinformers also inoculate their readers with critical citations of otherwise widely-trusted research. We reflect that mainstream science communicators' apparent reluctance to cite critically may stem from not wanting to attract attention, even critically, to counter-consensus findings. Though some research supports the benefits of prebunking (Ecker, Lewandowsky, and Chadwick, 2020), other findings have claimed that the "oxygen of amplification" may ultimately be a negative factor, and that counter-consensus sources are best left uncovered by mainstream media (Skurnik et al., 2005; Phillips, 2018). Those outside the scientific mainstream, however, need not worry about amplifying their opponents, as they may assume that most people will inevitably encounter the mainstream position. We recommend further research on the circumstances in which it may or may not be advantageous to negatively cite information that audience members may not have previously encountered.

Though we study Twitter here, it is likely the pattern we find persists across other forms of media. In fact, a quotation from one of the most-shared papers in the anti-mask consensus that we identified, MacIntyre et al. (MacIntyre et al., 2015), was cited in 2021 on what was at the time the most popular broadcast television program in the United States, Tucker Carlson Tonight, and in a 2022 legal brief successfully advocating for the abolition of federal mask mandates on public transportation (Hitlin, 2016; *Tucker Carlson Tonight* 2021). This is particularly unfortunate because Dr. MacIntyre herself has worked to support mask-wearing in many subsequent letters and peer-reviewed publications, and even specifically qualified the results of the 2015 paper in question (MacIntyre et al., 2020). While experts in public health have the ability to assess MacIntyre's 2015 paper holistically in light of her more recent work, non-experts, who encounter the 2015 paper in isolation on the journal's website or aggregators like PubMed, likely do not. This phenomenon has previously been called context collapse, where one group of users (non-experts) encounters

information meant for a different group (experts) and consequently lacks the resources to correctly reason about it (Marwick and Boyd, 2011). A key qualitative result from this work is how unprepared scientific publishers are for the implications of context collapse on their published archives.

To overcome context collapse, publishers could allow experts to append contextual information to papers that have been or are likely to be misused, and do so in formats that are difficult for non-experts to ignore, i.e. appending a large, conspicuous warning box rather than the small ones we find for these papers (e.g. in Klompas et al., 2020)). Sites which rehost scientific content, like PubMed and ResearchGate, must show a willingness to propagate such contextualizations, as the content in our study was often linked to these sites. Scientists themselves must show a willingness to create such contextualizations, and scientific institutions must create credit mechanisms that reward the creation of these contextualizations. Scientists' peers must understand and acknowledge that rather than a show of weakness on the part of the original authors' work, these additional comments represent valuable public service. Scientists, reviewers, and editors should also take care to craft their papers' abstracts with specific statements that are not so easily misquoted, as paper abstracts are sometimes the only information available for lay-readers without subscriptions to closed-access research. Such interventions can make the knowledge infrastructure of science more resilient to misuse by bad actors, while preserving the value of that infrastructure to expert scientists. We encourage additional research into what forms such interventions may take, and which are effective in defanging the work of counter-consensus science communicators (Smith, Nevarez, and Zhu, 2020).

The implication of this work that consensuses are socially constructed by non-expert mediators may cause familiar anxiety among scientists, as it can suggest that mainstream scientific consensus is non-objective and thus potentially just as untrustworthy (Kofman, 2018) as alternative consensuses. We take the perspective of Oreskes in her work *Why Trust Science*, and believe that the

selection of a certain consensus should depend on the diversity and strength of the vetting system that a consensus's authors practice (Oreskes, 2021). The architects of the anti-mask consensus found here are largely an informal coalition of conspiracy theorists, creators on alternative media platforms, and others who have previously shared false theories about climate change, the benefits of vaccination, the existence of mass shootings, HIV/AIDS, abortion, the reality of the moon landing, and a variety of other topics. Few have any scientific background in mask-effectiveness or even public health more generally, and similarly few have any formal accountability to the public. By contrast, the architects of the pro-mask consensus are a coalition of mostly journalists, scientists, and government agencies, all of whom have established different formal systems of peer review and accountability. While it is widely known that these institutions do not always produce reliable knowledge, especially for matters concerning marginalized people (TallBear, 2013; Hogarth, 2017; Dusenbery, 2018; Shuster, 2021), their systems for generating and vetting information are still stronger than those of a loose movement of unaffiliated conspiracy theorists and creators on alternative media platforms. As a community, scientists and publishers must work to build better, more trustworthy consensuses by strengthening both the popular perception — and the actual reality — of a science informed by rigorous peer review that is performed by and accountable to a diverse array of knowledgeable stakeholders (Fiske and Dupree, 2014). And they must understand that their work is filtered through the lens of science communicators, both friendly and unfriendly, and participate in the process of correction and communication that leads to their work being interpreted as intended.

We conclude with some comments on the construction of the datasets used in this study, which has been undoubtedly time-consuming, and how this process may be automated in the future. While automated positive/negative citation coding would allow analyses of larger controversies over longer periods of time, our experience manually coding citations suggests that automated processes will face difficulties. We find that scientists in archival work do not criticize papers the way people criticize, for example, restaurants or movies, and will often use non-generalizable

science- and field-specific language. Criticisms can be circumspect due to admirable comity norms in science, and their meaning may be diffused over several sentences. Outside of automatic coding, simply gathering the citations in a consistent manner is likely to be quite difficult. Websites communicating important scientific information can change their content almost daily, which we observed particularly with the CDC's websites, and this complicates attempts to understand the state of perceived scientific consensus at any given moment. Automatically collecting citations from websites with very different layouts and citation formats will also inevitably provide challenges, especially given how many of these websites do not hyperlink (or even consistently format) their citations. The popularity of video-based science communication, which combines audio content with images and text screenshots of scientific papers sometimes without titles, further increases the difficulty of completing an automatic census of communicated scientific literature as actually experienced by the media-consuming public. While we encourage research into how such data collection may be automated, we also encourage researchers to continue to engage in similar hand-coding experiments in scientific communication in order to surface such challenges.

7.6 Materials and Methods

7.6.1 Web of Science Dataset

To understand how scientists were developing consensus around masks in the wake of the COVID-19 pandemic, we used the Web of Science paper database to collect mask-related papers in approximately the first year of the pandemic. Specifically, we collect 3,593 articles from the Web of Science published between March 28, 2020 and April 4, 2021 containing in their title the following: "mask," "personal protective equipment," "ppe", and "n95." These choices were informed by familiarizing ourselves with popular papers referencing masks, and the terms most frequently used to describe masks. We filter this dataset to the 82 papers that have at least 50 citations as of November 18, 2021, representing the top 2% of papers retrieved. We choose this threshold to limit the labor needed for subsequent manual coding. We exclude two papers unrelated to masking

during the COVID-19 pandemic, and three papers with no collectable citations, leaving a dataset of 77 papers. These papers are listed in the supplemental information.

Most papers are peer-reviewed experiments and literature reviews; others are non-peer-reviewed editorials. Most directly pertain to measuring the effectiveness of masks for preventing the spread of COVID-19, but some pertain to adjacent issues such as the effectiveness of re-using masks or mitigating the waste products created from COVID-19-related mask production.

For the 77 papers, we retrieve all listed outgoing citations within the papers that are archived by the Web of Science, using the Digital Object Identifier (DOI). While this does not include all citations, our objective is to replicate those data curation practices typically performed in large-scale bibliometric studies of science. Even including citations without DOIs, the Web of Science does not include every citation that these articles make, especially when citations are made to work published outside of academic journals. There are rare instances in which the Web of Science consistently misrouted a given paper's citation to the wrong paper, and we attempted to correct these when identified. This process results in 2,426 citations, with the most cited paper as Leung et al. (Leung et al., 2020) referenced in 29 papers.

7.6.2 Twitter Science Dataset

Most social media posts do not necessarily constitute scientific debate, which poses a challenge from a data collection perspective. A naive collection of Twitter posts containing the word "mask" during our study period surfaces mostly content related to sensational stories about masks. Such links rarely contain scientific content, and instead document, for example, an altercation in a gas station started over a request to wear a mask. To discover users' content about the effectiveness of masks, we instead specifically sample reply threads between users discussing masks — often in disagreement with one other. These arguments stimulate the trading of evidence between users, in the form of external URLs, which then become the object of analysis for the Twitter Science Dataset used here.

We choose as our forum for arguments posts made in reply to the set of all Twitter accounts for United States governors (and the mayor of the District of Columbia) as of March 28, 2020. These political leaders are often tasked with announcing and explaining public health regulations that require wearing masks during the COVID-19 pandemic, and their posts about masks tend to generate fierce arguments between pro- and anti-mask users. From March 28, 2020 to April 4, 2021, we collect all hierarchical reply trees to these politicians which contain the word "mask" at any point in the tree. The number of posts in all reply trees totalled 37.8M, and the subset of trees containing an instance of "mask" totalled 5.1M posts (13% of the total). From these posts, we extracted all posts which shared a URL for a total of 158K unique URLs shared across 351K posts. We collected Twitter posts using the Twitter Streaming API, and independently unwound redirected links to retrieve the final destination of original URLs. Some URLs could not be retrieved due to redirect URLs no longer being operational.

From this dataset of URLs, we first remove marketing tags and other character strings appended to the end of URLs that do not affect the final landing page, and then deduplicate the results. This reduces the URL dataset size by 10% to 142K unique URLs. To find a tractable sample for manual coding, we subsample the 194 URLs which were shared more than 100 times, which collectively are responsible for 15% of all URLs shared. Two researchers then perform a qualitative coding process on the linked content to identify which are explicitly concerned with the effectiveness (or ineffectiveness) of mask-wearing for COVID-19. This restricted set of links comprises 98 unique URLs. We combined unique URLs which rehosted identical content on different websites, for a final set of 81 unique pieces of science communication. Many URLs had either been modified or, in the case of social media posts, deleted by the time of coding. To address these gaps, we use the Internet Archive (Kahle, 1997) to access the most recent non-redirected version as of April 4, 2021. In only one case could a URL, linking to a removed Twitter post, not be retrieved via this method. We note that 8/10 of the most shared unique URLs specifically concerned themselves with a scientific analysis of the effectiveness of masks, with the remaining non-scientific links

corresponding to a post by former United States president Donald Trump and a petition to recall California Governor Gavin Newsom.

From this dataset of 81 URLs, two researchers manually recorded every reference in the main text of these URLs. These references could be represented as embedded URLs, appended as formal academic citations, screenshots of text, or explicitly referenced by article name. In the case of URLs leading to videos, researchers watched the videos and recorded verbal citations and citations displayed in text on the video itself, a frequent behavior. All citations to external sources were recorded, including material published in scientific journals as well as sources such as videos, scientific journalism, and other websites. We coded the most recent version of the website as of April 4, 2021, and only content on the site which directly pertain to masks and the article in question. This resulted in 1,496 citations, with the most-cited article (as in the Web of Science dataset) being Leung et al., cited by 19 unique URLs (Leung et al., 2020).

7.6.3 Annotating citation and website valence

In both the Web of Science and the Twitter Citations datasets, the lead researcher manually annotated citations for valence, specifically noting negative citations. We define two criteria for a negative citation:

- Criticizing the methods, conclusions, or statistical power of a cited study (e.g. "Many studies in these reviews were underpowered, and most failed to measure adherence..." (Javid, Weekes, and Matheson, 2020))
- Arguing that a cited study has been popularly misinterpreted by peers or the public, either in general or for the specific case in context (e.g. "One of the most frequently mentioned, but misinterpreted, papers evaluating cloth masks..." (Howard et al., 2020)).

We do not label citations negatively when authors raise limitations or areas for improvement in otherwise trustworthy work, or in the case where cited findings conflict with the findings of the

present paper but are not themselves criticized. Our intended definition for negative citations is that, in effect, the citing article is asking readers to at least partially disregard the conclusions of a cited paper.

For our analysis of negative citation rates in Twitter science communicators, we also labeled websites in our dataset as either having a stance generally critical of masks as intervention for COVID-19, or generally positive or neutral. We considered websites to have an anti-mask orientation if they either explicitly advocated against the wearing of masks or primarily reported scientific findings supporting the ineffectiveness of masks. Anti-mask websites which had been publicly corrected and modified by the time of our analysis were considered to have neutral stance. We labeled 40 websites as critical of masks, and 41 as positive or neutral.

7.6.4 Network Construction

Co-citation graphs are essentially undirected unipartite projections of directed bipartite graphs, where an edge between two nodes signifies that two papers are cited together in at least one citing article (Small, 1973). Projecting a bipartite network into a unipartite network is never trivial, and choices about how to represent edges and edge weights can significantly affect resulting analyses, particularly clusterings (Cann, Weaver, and Williams, 2020). These choices are made even more complex in the signed case we introduce here. We perceive three principal researcher choices with respect to these projections: 1) the choice to analyze a unipartite projection rather than the original bipartite network, 2) the scheme for generating edges and edge weights for the unipartite projection, and 3) the scheme for filtering edges such that communities of interest are preserved.

The first is the decision to project a bipartite graph, especially when the original graph could be considered unipartite given that some sources cite other sources in the citing set. Our choice to treat "citing" and "cited" entities as separate and thus bipartite stems from our understanding

of this data as representing the curatorial choices of a source's authors, rather than a deep engagement with a source's content. Indeed, we find that many anti-mask communicators misleadingly cite research that actually endorses masks and itself cites extensively pro-mask literature. In directed unipartite form, this would lead to anti-mask researchers seemingly endorsing a wide variety of pro-mask literature via one transitive connection.

The second choice is how to define edges in a signed projected graph. We follow a method described in (Derr et al., 2019), where positive edges are considered to be (+1) and negative edges are considered to be (-1), and the weight contribution to an edge cited by two nodes being the product of their cited edges (1*1=1, -1*-1=1, 1*-1=-1). Nodes are either together within trustworthy science (1*1), together within untrustworthy pseudoscience (-1*-1), or apart from each other on either side of the boundary (1*-1). In subsequent analysis, we analyze two such projections: one which contains all of the summed positive edge contributions (nodes cited with the same valence) and one which contains all of the negative edge contributions (nodes cited with opposite valence) over the same set of nodes. Splitting these projections, rather than summing them into one graph, aids in cluster detection via methods which optimize partitionings between multigraphs, as described in the following section.

The third choice is how to weight and filter the edges to only those relevant to the clusters at hand. As described in (Cann, Weaver, and Williams, 2020), projections tend to generate local structures which are highly amenable to clustering via modularity based methods, because each citing node creates a densely-connected subgraph of all the nodes it cites in the projection. Accordingly, modularity-based methods can easily return clusters that simply represent "all the papers cited by a single source." While weighting methods have been proposed that mitigate the cluster-making ability of single sources, such as hyperbolic weight projection, we find that their application to signed citation networks nontrivial, and not necessarily well-suited to the relatively normal degree-distribution found in our data. We instead opt for a simple filtering mechanism: connections between two papers are not considered unless they have been cited together in at

least three sources in our datasets. This eliminates the cluster-making ability of single authors, and single authors who create two sources and use similar citation lists. Our data does not feature authors with more than two sources, making stricter filters unnecessary. We note that this is not a simple filter on edge weights less than 3 on both subgraphs, as two sources could be cited positively together in two sources (+2) and negatively together in one (-1), and these edges would still be preserved in their respective projections because they represent co-citation in three sources.

7.6.5 Clustering and Visualization

Our primary clustering algorithm is based on the modularity statistic (Newman and Girvan, 2004), and optimizes modularity using a multigraph method derived from the Leiden algorithm (Traag, Waltman, and Eck, 2019) as implemented in the Python package (Traag and Bruggeman, 2009). Specifically, this optimization method attempts to maximize modularity on the graph representing positive ties while simultaneously minimizing modularity on the graph representing negative ties.

Historically, co-citation analyses use large, automatically-collected datasets with hundreds to millions of nodes and edges. Our case is unique in that, due to hand-coding, our network contains only hundreds of nodes and thousands of connections. The small size of our dataset increases the influence of noise in our data, and correspondingly the level of uncertainty in clusterings and network representations. To address this challenge, we use a method akin to soft-clustering, where individual nodes are assigned cluster probabilities rather than strict assignments, and inspired by the clustering bootstrapping procedure proposed in (Rosvall and Bergstrom, 2010). Specifically, we resample each weight in the projected graphs according to a normal distribution with mean of the given weight, and variance equal to the variance of all weights in the graph. We perform such replications and cluster the results 1,000 times for each graph, permuting the order of nodes provided to algorithm each time. We then create a derived representation of these graphs where each node represents a cited paper, and links between nodes are weighted by the percentage (0-1.0) of times these nodes were clustered in the same graph. We then apply clustering using an unsigned

modularity method to this derived graph. We visualize these representations in Figure 1, with a filter such that nodes that are clustered less than 50% of the time are removed.

We visualize graphs using the ForceAtlas2 algorithm in the network visualization software Gephi (Bastian, Heymann, and Jacomy, 2009; Jacomy et al., 2014). Nodes in Figure 1 are sized by the summed number of citations the papers citing them have received, to illustrate their importance in scientific discourse around masks in the first year of the pandemic. Nodes in Figure 1 are similarly sized by the summed number of times the websites citing them have been shared on Twitter, for a parallel measure of importance.

7.7 Acknowledgements

We thank Joseph S. Schafer, Jueqi Liu, Mariam Mayanja, Monica Ionescu, Maya Sioson, and Paul Lockaby for their advice and support on this work.

7.8 Funding

This work has been supported by the University of Washington's Center for an Informed Public, the John S. and James L. Knight Foundation (G-2019-58788), the William and Flora Hewlett Foundation (2019-9221), and the University of Washington Population Health Initiative Pilot Grants. We also are grateful for support from the National Science Foundation (grants 2027792 and 1749815).

7.9 Author Contributions

Conceptualization: AB, SN

Methodology: AB, JDW, ESS

Investigation: AB, SN

Visualization: AB

Supervision: KS, JDW, ESS

Writing—original draft: AB

Writing—review & editing: AB, SN, KS, JDW, ESS

7.10 Competing Interests

The authors declare that they have no competing interests. The full coded citation data for this paper can be downloaded here:

7.11 Data and Materials Availability

The code for reproducing this paper, Twitter URL dataset, Web of Science citation dataset, annotation datasets, and network visualization data can be accessed online via Zenodo: https://zenodo.org/record/803
All data needed to evaluate the conclusions of the paper are present in the paper and/or the supplementary materials.

7.12 Appendix

The 82 highly-cited masks papers we use in the "Web of Science" dataset are listed as follows:

S1. S. Feng, C. Shen, N. Xia, W. Song, M. Fan, B. J. Cowling, Rational use of face masks in the COVID-19 pandemic. The Lancet Respiratory Medicine. 8, 434–436 (2020).

S2. A. Schwartz, M. Stiegel, N. Greeson, A. Vogel, W. Thomann, M. Brown, G. D. Sempowski, T. S. Alderman, J. P. Condreay, J. Burch, C. Wolfe, B. Smith, S. Lewis, Decontamination and Reuse of N95 Respirators with Hydrogen Peroxide Vapor to Address Worldwide Personal Protective Equipment Shortages During the SARS-CoV-2 (COVID-19) Pandemic. Applied Biosafety. 25, 67–70 (2020).

S3. M. L. Ranney, V. Griffeth, A. K. Jha, Critical Supply Shortages — The Need for Ventilators and Personal Protective Equipment during the Covid-19 Pandemic. New England Journal of Medicine. 382, e41 (2020).

S4. J. Xiao, E. Y. C. Shiu, H. Gao, J. Y. Wong, M. W. Fong, S. Ryu, B. J. Cowling, Nonpharmaceutical Measures for Pandemic Influenza in Nonhealthcare Settings—Personal Protective and Environmental Measures. Emerging Infectious Diseases. 26, 967–975 (2020).

S5. X. Wang, Z. Pan, Z. Cheng, Association between 2019-nCoV transmission and N95 respirator use. Journal of Hospital Infection. 105, 104–105 (2020).

S6. A. Tabah, M. Ramanan, K. B. Laupland, N. Buetti, A. Cortegiani, J. Mellinghoff, A. C. Morris, L. Camporota, N. Zappella, M. Elhadi, P. Povoa, K. Amrein, G. Vidal, L. Derde, M. Bassetti, G. Francois, N. S. yan kai, J. J. D. Waele, Personal protective equipment and intensive care unit healthcare worker safety in the COVID-19 era (PPE-SAFE): An international survey. Journal of Critical Care. 59, 70–75 (2020).

S7. E. N. Perencevich, D. J. Diekema, M. B. Edmond, Moving Personal Protective Equipment Into the Community. JAMA. 323, 2252 (2020).

S8. B. E. McGarry, D. C. Grabowski, M. L. Barnett, Severe Staffing And Personal Protective Equipment Shortages Faced By Nursing Homes During The COVID-19 Pandemic. Health Affairs. 39, 1812–1821 (2020).

S9. Y. Long, T. Hu, L. Liu, R. Chen, Q. Guo, L. Yang, Y. Cheng, J. Huang, L. Du, Effectiveness of N95 respirators versus surgical masks against influenza: A systematic review and meta-analysis. Journal of Evidence-Based Medicine. 13, 93–101 (2020).

S10. E. Livingston, A. Desai, M. Berkwits, Sourcing Personal Protective Equipment During the COVID-19 Pandemic. JAMA. 323, 1912 (2020). S11. L. Liao, W. Xiao, M. Zhao, X. Yu, H. Wang, Q. Wang, S. Chu, Y. Cui, Can N95 Respirators Be Reused after Disinfection? How Many Times? ACS Nano. 14, 6348–6356 (2020).

S12. Z. M. Jessop, T. D. Dobbs, S. R. Ali, E. Combellack, R. Clancy, N. Ibrahim, T. H. Jovic, A. J. Kaur, A. Nijran, T. B. O'Neill, I. S. Whitaker, Personal protective equipment for surgeons during COVID-19 pandemic: systematic review of availability, usage and rationing. British Journal of Surgery. 107, 1262–1280 (2020).

S13. M. T. Hirschmann, A. Hart, J. Henckel, P. Sadoghi, R. Seil, C. Mouton, COVID-19 coronavirus: recommended personal protective equipment for the orthopaedic and trauma surgeon. Knee Surgery, Sports Traumatology, Arthroscopy. 28, 1690–1698 (2020).

S14. T. M. Cook, Personal protective equipment during the coronavirus disease (COVID) 2019 pandemic – a narrative review. Anaesthesia. 75, 920–927 (2020). S15. J. Cohen, Y. van der M.

Rodgers, Contributing factors to personal protective equipment shortages during the COVID-19 pandemic. Preventive Medicine. 141, 106263 (2020).

S16. C. Y. Benítez, A. Güemes, J. Aranda, M. Ribeiro, P. Ottolino, S. D. Saverio, H. Alexandrino, L. Ponchietti, J. L. Blas, J. P. Ramos, E. Rangelova, M. Muñoz, C. Y. and, Impact of Personal Protective Equipment on Surgical Performance During the COVID-19 Pandemic. World Journal of Surgery. 44, 2842–2847 (2020).

S17. J. J. Bartoszko, M. A. M. Farooqi, W. Alhazzani, M. Loeb, Medical masks vs N95 respirators for preventing COVID-19 in healthcare workers: A systematic review and meta-analysis of randomized trials. Influenza and Other Respiratory Viruses. 14, 365–373 (2020).

S18. Y. Wang, H. Tian, L. Zhang, M. Zhang, D. Guo, W. Wu, X. Zhang, G. L. Kan, L. Jia, D. Huo, B. Liu, X. Wang, Y. Sun, Q. Wang, P. Yang, C. R. MacIntyre, Reduction of secondary transmission of SARS-CoV-2 in households by face mask use, disinfection and social distancing: a cohort study in Beijing, China. BMJ Global Health. 5, e002794 (2020).

S19. C. Wang, A. Chudzicka-Czupala, D. Grabowski, R. Pan, K. Adamus, X. Wan, M. Hetnał, Y. Tan, A. Olszewska-Guizzo, L. Xu, R. S. McIntyre, J. Quek, R. Ho, C. Ho, The Association Between Physical and Mental Health and Face Mask Use During the COVID-19 Pandemic: A Comparison of Two Countries With Different Views and Practices. Frontiers in Psychiatry. 11 (2020), doi:10.3389/fpsyt.2020.569981.

S20. R. Sommerstein, and C. A. Fux, D. Vuichard-Gysin, M. Abbas, J. Marschall, C. Balmelli, N. Troillet, S. Harbarth, M. Schlegel, A. Widmer, Risk of SARS-CoV-2 transmission by aerosols, the rational use of masks, and protection of healthcare workers from COVID-19. Antimicrobial Resistance & Infection Control. 9 (2020), doi:10.1186/s13756-020-00763-0.

S21. J. C. Rubio-Romero, M. del C. Pardo-Ferreira, J. A. Torrecilla-García, S. Calero-Castro, Disposable masks: Disinfection and sterilization for reuse, and non-certified manufacturing, in the face of shortages during the COVID-19 pandemic. Safety Science. 129, 104830 (2020).

S22. K. A. Prather, C. C. Wang, R. T. Schooley, Reducing transmission of SARS-CoV-2. Science. 368, 1422–1424 (2020).

S23. S. Pfattheicher, L. Nockur, R. Böhm, C. Sassenrath, M. B. Petersen, The Emotional Path to Action: Empathy Promotes Physical Distancing and Wearing of Face Masks During the COVID-19 Pandemic. Psychological Science. 31, 1363–1373 (2020).

S24. R. Mittal, C. Meneveau, W. Wu, A mathematical framework for estimating risk of airborne transmission of COVID-19 with application to face mask use and social distancing. Physics of Fluids. 32, 101903 (2020).

S25. C. R. MacIntyre, A. A. Chughtai, A rapid systematic review of the efficacy of face masks and respirators against coronaviruses and other respiratory transmissible viruses for the community, healthcare workers and sick patients. International Journal of Nursing Studies. 108, 103629 (2020).

S26. W. Lyu, G. L. Wehby, Community Use Of Face Masks And COVID-19: Evidence From A Natural Experiment Of State Mandates In The US. Health Affairs. 39, 1419–1425 (2020).

S27. M. Loey, G. Manogaran, M. H. N. Taha, N. E. M. Khalifa, A hybrid deep transfer learning model with machine learning methods for face mask detection in the era of the COVID-19 pandemic. Measurement. 167, 108288 (2021).

S28. S. L. Lockhart, L. V. Duggan, R. S. Wax, S. Saad, H. P. Grocott, Personal protective equipment (PPE) for both anesthesiologists and other airway managers: principles and practice during the COVID-19 pandemic. Canadian Journal of Anesthesia/Journal canadien d'anesthésie. 67, 1005–1015 (2020).

S29. M. Gandhi, C. Beyrer, E. Goosby, Masks Do More Than Protect Others During COVID-19: Reducing the Inoculum of SARS-CoV-2 to Protect the Wearer. Journal of General Internal Medicine. 35, 3063–3066 (2020).

S30. S. Fikenzer, T. Uhe, D. Lavall, U. Rudolph, R. Falz, M. Busse, P. Hepp, U. Laufs, Effects of surgical and FFP2/N95 face masks on cardiopulmonary exercise capacity. Clinical Research in Cardiology. 109, 1522–1530 (2020).

S31. D. K. Chu, E. A. Akl, S. Duda, K. Solo, S. Yaacoub, H. J. Schünemann, D. K. Chu, E. A. Akl, A. El-harakeh, A. Bognanni, T. Lotfi, M. Loeb, A. Hajizadeh, A. Bak, A. Izcovich, C. A. Cuello-Garcia, C. Chen, D. J. Harris, E. Borowiack, F. Chamseddine, F. Schünemann, G. P. Morgano, G. E. U. M. Schünemann, G. Chen, H. Zhao, I. Neumann, J. Chan, J. Khabsa, L. Hneiny, L. Harrison, M. Smith, N. Rizk, P. G. Rossi, P. AbiHanna, R. El-khoury, R. Stalteri, T. Baldeh, T. Piggott, Y. Zhang, Z. Saad, A. Khamis, M. Reinap, S. Duda, K. Solo, S. Yaacoub, H. J. Schünemann, Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: a systematic review and meta-analysis. The Lancet. 395, 1973–1987 (2020).

S32. R. Chou, T. Dana, R. Jungbauer, C. Weeks, M. S. McDonagh, Masks for Prevention of Respiratory Virus Infections, Including SARS-CoV-2, in Health Care and Community Settings. Annals of Internal Medicine. 173, 542–555 (2020).

S33. V. C.-C. Cheng, S.-C. Wong, V. W.-M. Chuang, S. Y.-C. So, J. H.-K. Chen, S. Sridhar, K. K.-W. To, J. F.-W. Chan, I. F.-N. Hung, P.-L. Ho, K.-Y. Yuen, The role of community-wide wearing of face mask for control of coronavirus disease 2019 (COVID-19) epidemic due to SARS-CoV-2. Journal of Infection. 81, 107–114 (2020).

S34. X. Chen, L. Ran, Q. Liu, Q. Hu, X. Du, X. Tan, Hand Hygiene, Mask-Wearing Behaviors and Its Associated Factors during the COVID-19 Epidemic: A Cross-Sectional Study among Primary School Students in Wuhan, China. International Journal of Environmental Research and Public Health. 17, 2893 (2020).

S35. J. F.-W. Chan, S. Yuan, A. J. Zhang, V. K.-M. Poon, C. C.-S. Chan, A. C.-Y. Lee, Z. Fan, C. Li, R. Liang, J. Cao, K. Tang, C. Luo, V. C.-C. Cheng, J.-P. Cai, H. Chu, K.-H. Chan, K. K.-W. To, S. Sridhar, K.-Y. Yuen, Surgical Mask Partition Reduces the Risk of Noncontact Transmission in a Golden Syrian Hamster Model for Coronavirus Disease 2019 (COVID-19). Clinical Infectious Diseases. 71, 2139–2149 (2020).

S36. C. D. Zangmeister, J. G. Radney, E. P. Vicenzi, J. L. Weaver, Filtration Efficiencies of Nanoscale Aerosol by Cloth Mask Materials Used to Slow the Spread of SARS-CoV-2. ACS Nano. 14, 9188–9200 (2020).

S37. C. J. Worby, H.-H. Chang, Face mask use in the general population and optimal resource allocation during the COVID-19 pandemic. Nature Communications. 11 (2020), doi:10.1038/s41467-020-17922-x.

S38. X. Wang, E. G. Ferro, G. Zhou, D. Hashimoto, D. L. Bhatt, Association Between Universal Masking in a Health Care System and SARS-CoV-2 Positivity Among Health Care Workers. JAMA. 324, 703 (2020).

S39. S. Ullah, A. Ullah, J. Lee, Y. Jeong, M. Hashmi, C. Zhu, K. I. Joo, H. J. Cha, I. S. Kim, Reusability Comparison of Melt-Blown vs Nanofiber Face Mask Filters for Use in the Coronavirus Pandemic. ACS Applied Nano Materials. 3, 7231–7241 (2020).

S40. A. Tcharkhtchi, N. Abbasnezhad, M. Z. Seydani, N. Zirak, S. Farzaneh, M. Shirinbayan, An overview of filtration efficiency through the masks: Mechanisms of the aerosols penetration. Bioactive Materials. 6, 106–122 (2021).

S41. N. J. Rowan, J. G. Laffey, Challenges and solutions for addressing critical shortage of supply chain for personal and protective equipment (PPE) arising from Coronavirus disease (COVID19) pandemic – Case study from the Republic of Ireland. Science of The Total Environment. 725, 138532 (2020).

S42. Q.-X. Ma, H. Shan, H.-L. Zhang, G.-M. Li, R.-M. Yang, J.-M. Chen, Potential utilities of mask-wearing and instant hand hygiene for fighting SARS-CoV-2. Journal of Medical Virology. 92, 1567–1571 (2020).

S43. M. Liang, L. Gao, C. Cheng, Q. Zhou, J. P. Uy, K. Heiner, C. Sun, Efficacy of face mask in preventing respiratory virus transmission: A systematic review and meta-analysis. Travel Medicine and Infectious Disease. 36, 101751 (2020).

S44. M. Ippolito, F. Vitale, G. Accurso, P. Iozzo, C. Gregoretti, A. Giarratano, A. Cortegiani, Medical masks and Respirators for the Protection of Healthcare Workers from SARS-CoV-2 and other viruses. Pulmonology. 26, 204–212 (2020).

S45. M. H. Haischer, R. Beilfuss, M. R. Hart, L. Opielinski, D. Wrucke, G. Zirgaitis, T. D. Uhrich, S. K. Hunter, Who is wearing a mask? Gender-, age-, and location-related differences during the COVID-19 pandemic. PLOS ONE. 15, e0240785 (2020).

S46. B. Ghanbari, A. Atangana, A new application of fractional Atangana–Baleanu derivatives: Designing ABC-fractional masks in image processing. Physica A: Statistical Mechanics and its Applications. 542, 123516 (2020).

S47. S. Esposito, N. Principi, C. C. Leung, G. B. Migliori, Universal use of face masks for success against COVID-19: evidence and implications for prevention policies. European Respiratory Journal. 55, 2001260 (2020).

S48. N. El-Atab, N. Qaiser, H. Badghaish, S. F. Shaikh, M. M. Hussain, Flexible Nanoporous Template for the Design and Development of Reusable Anti-COVID-19 Hydrophobic Face Masks. ACS Nano. 14, 7659–7665 (2020).

S49. S. Bae, M.-C. Kim, J. Y. Kim, H.-H. Cha, J. S. Lim, J. Jung, M.-J. Kim, D. K. Oh, M.-K. Lee, S.-H. Choi, M. Sung, S.-B. Hong, J.-W. Chung, S.-H. Kim, Effectiveness of Surgical and Cotton Masks in Blocking SARS–CoV-2: A Controlled Comparison in 4 Patients. Annals of Internal Medicine. 173, W22–W23 (2020).

S50. S. Asadi, C. D. Cappa, S. Barreda, A. S. Wexler, N. M. Bouvier, W. D. Ristenpart, Efficacy of masks and face coverings in controlling outward aerosol particle emission from expiratory activities. Scientific Reports. 10 (2020), doi:10.1038/s41598-020-72798-7.

S51. T. A. Aragaw, Surgical face masks as a potential source for microplastic pollution in the COVID-19 scenario. Marine Pollution Bulletin. 159, 111517 (2020).

S52. P. C. Addo, F. Jiaming, N. B. Kulbo, L. Liangqiang, COVID-19: fear appeal favoring purchase behavior towards personal protective equipment. The Service Industries Journal. 40, 471–490 (2020).

S53. H. Zhong, Z. Zhu, J. Lin, C. F. Cheung, V. L. Lu, F. Yan, C.-Y. Chan, G. Li, Reusable and Recyclable Graphene Masks with Outstanding Superhydrophobic and Photothermal Performances. ACS Nano. 14, 6213–6221 (2020).

S54. S. Verma, M. Dhanak, J. Frankenfield, Visualizing droplet dispersal for face shields and masks with exhalation valves. Physics of Fluids. 32, 091701 (2020).

S55. S. Verma, M. Dhanak, J. Frankenfield, Visualizing the effectiveness of face masks in obstructing respiratory jets. Physics of Fluids. 32, 061708 (2020).

S56. H. Ueki, Y. Furusawa, K. Iwatsuki-Horimoto, M. Imai, H. Kabata, H. Nishimura, Y. Kawaoka, Effectiveness of Face Masks in Preventing Airborne Transmission of SARS-CoV-2. mSphere. 5 (2020), doi:10.1128/msphere.00637-20.

S57. A. Pezzini, A. Padovani, Lifting the mask on neurological manifestations of COVID-19. Nature Reviews Neurology. 16, 636–644 (2020).

S58. L. J. Mady, M. W. Kubik, K. Baddour, C. H. Snyderman, N. R. Rowan, Consideration of povidone-iodine as a public health intervention for COVID-19: Utilization as "Personal Protective Equipment" for frontline providers exposed in high-risk head and neck and skull base oncology care. Oral Oncology. 105, 104724 (2020).

S59. X. Liu, S. Zhang, COVID-19: Face masks and human-to-human transmission. Influenza and Other Respiratory Viruses. 14, 472–473 (2020).

S60. N. H. L. Leung, D. K. W. Chu, E. Y. C. Shiu, K.-H. Chan, J. J. McDevitt, B. J. P. Hau, H.-L. Yen, Y. Li, D. K. M. Ip, J. S. M. Peiris, W.-H. Seto, G. M. Leung, D. K. Milton, B. J. Cowling, Respiratory virus shedding in exhaled breath and efficacy of face masks. Nature Medicine. 26, 676–680 (2020).

S61. C. C. Leung, T. H. Lam, K. K. Cheng, Mass masking in the COVID-19 epidemic: people need guidance. The Lancet. 395, 945 (2020).

S62. A. Konda, A. Prakash, G. A. Moss, M. Schmoldt, G. D. Grant, S. Guha, Aerosol Filtration Efficiency of Common Fabrics Used in Respiratory Cloth Masks. ACS Nano. 14, 6339–6347 (2020).

S63. J. Howard, A. Huang, Z. Li, Z. Tufekci, V. Zdimal, H.-M. van der Westhuizen, A. von Delft, A. Price, L. Fridman, L.-H. Tang, V. Tang, G. L. Watson, C. E. Bax, R. Shaikh, F. Questier, D. Hernandez, L. F. Chu, C. M. Ramirez, A. W. Rimoin, An evidence review of face masks against COVID-19. Proceedings of the National Academy of Sciences. 118 (2021), doi:10.1073/pnas.2014564118.

S64. T. Greenhalgh, M. B. Schmid, T. Czypionka, D. Bassler, L. Gruer, Face masks for the public during the covid-19 crisis. BMJ, m1435 (2020).

S65. M. Gandhi, G. W. Rutherford, Facial Masking for Covid-19 — Potential for "Variolation" as We Await a Vaccine. New England Journal of Medicine. 383, e101 (2020).

S66. E. P. Fischer, M. C. Fischer, D. Grass, I. Henrion, W. S. Warren, E. Westman, Low-cost measurement of face mask efficacy for filtering expelled droplets during speech. Science Advances. 6 (2020), doi:10.1126/sciadv.abd3083.

S67. O. O. Fadare, E. D. Okoffo, Covid-19 face masks: A potential source of microplastic fibers in the environment. Science of The Total Environment. 737, 140279 (2020).

S68. M. H. Chua, W. Cheng, S. S. Goh, J. Kong, B. Li, J. Y. C. Lim, L. Mao, S. Wang, K. Xue, L. Yang, E. Ye, K. Zhang, W. C. D. Cheong, B. H. Tan, Z. Li, B. H. Tan, X. J. Loh, Face Masks in the New COVID-19 Normal: Materials, Testing, and Perspectives. Research. 2020, 1–40 (2020).

S69. H. Bundgaard, J. S. Bundgaard, D. E. T. Raaschou-Pedersen, C. von Buchwald, T. Todsen, J. B. Norsk, M. M. Pries-Heje, C. R. Vissing, P. B. Nielsen, U. C. Winsløw, K. Fogh, R. Hasselbalch, J. H. Kristensen, A. Ringgaard, M. P. Andersen, N. B. Goecke, R. Trebbien, K. Skovgaard, T. Benfield, H. Ullum, C. Torp-Pedersen, K. Iversen, Effectiveness of Adding a Mask Recommendation to Other Public Health Measures to Prevent SARS-CoV-2 Infection in Danish Mask Wearers. Annals of Internal Medicine. 174, 335–343 (2021).

S70. J. T. Brooks, J. C. Butler, R. R. Redfield, Universal Masking to Prevent SARS-CoV-2 Transmission—The Time Is Now. JAMA. 324, 635 (2020).

S71. G. R. J. Swennen, L. Pottel, P. E. Haers, Custom-made 3D-printed face masks in case of pandemic crisis situations with a lack of commercially available FFP2/3 masks. International Journal of Oral and Maxillofacial Surgery. 49, 673–677 (2020).

S72. A. L. P. Silva, J. C. Prata, T. R. Walker, D. Campos, A. C. Duarte, A. M. V. M. Soares, D. Barcelò, T. Rocha-Santos, Rethinking and optimising plastic waste management under COVID-19 pandemic: Policy solutions based on redesign and reduction of single-use plastics and personal protective equipment. Science of The Total Environment. 742, 140565 (2020).

S73. S. W. X. Ong, Y. K. Tan, P. Y. Chia, T. H. Lee, O. T. Ng, M. S. Y. Wong, K. Marimuthu, Air, Surface Environmental, and Personal Protective Equipment Contamination by Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2) From a Symptomatic Patient. JAMA. 323, 1610 (2020).

S74. P. Mick, R. Murphy, Aerosol-generating otolaryngology procedures and the need for enhanced PPE during the COVID-19 pandemic: a literature review. Journal of Otolaryngology - Head & Neck Surgery. 49 (2020), doi:10.1186/s40463-020-00424-7.

S75. T. Li, Y. Liu, M. Li, X. Qian, S. Y. Dai, Mask or no mask for COVID-19: A public health and market study. PLOS ONE. 15, e0237691 (2020).

S76. M. Klompas, C. A. Morris, J. Sinclair, M. Pearson, E. S. Shenoy, Universal Masking in Hospitals in the Covid-19 Era. New England Journal of Medicine. 382, e63 (2020).

S77. J. J. Klemeš, Y. V. Fan, P. Jiang, The energy and environmental footprints of COVID-19 fighting measures – PPE, disinfection, supply chains. Energy. 211, 118701 (2020).

S78. B. Javid, M. P. Weekes, N. J. Matheson, Covid-19: should the public wear face masks? BMJ, m1442 (2020).

S79. T. Dbouk, D. Drikakis, On respiratory droplets and face masks. Physics of Fluids. 32, 063303 (2020).

S80. J. J. Y. Ong, C. Bharatendu, Y. Goh, J. Z. Y. Tang, K. W. X. Sooi, Y. L. Tan, B. Y. Q. Tan, H.-L. Teoh, S. T. Ong, D. M. Allen, V. K. Sharma, Headaches Associated With Personal Protective Equipment – A Cross-Sectional Study Among Frontline Healthcare Workers During COVID-19. Headache: The Journal of Head and Face Pain. 60, 864–877 (2020).

S81. M. Liu, S.-Z. Cheng, K.-W. Xu, Y. Yang, Q.-T. Zhu, H. Zhang, D.-Y. Yang, S.-Y. Cheng, H. Xiao, J.-W. Wang, H.-R. Yao, Y.-T. Cong, Y.-Q. Zhou, S. Peng, M. Kuang, F.-F. Hou, K. K. Cheng, H.-P. Xiao, Use of personal protective equipment against coronavirus disease 2019 by healthcare professionals in Wuhan, China: cross sectional study. BMJ, m2195 (2020).

S82. R. J. Fischer, D. H. Morris, N. van Doremalen, S. Sarchette, M. J. Matson, T. Bushmaker, C. K. Yinda, S. N. Seifert, A. Gamble, B. N. Williamson, S. D. Judson, E. de Wit, J. O. Lloyd-Smith, V. J. Munster, Effectiveness of N95 Respirator Decontamination and Reuse against SARS-CoV-2 Virus. Emerging Infectious Diseases. 26, 2253–2255 (2020).

Chapter 8

Discussion and Conclusion

In this chapter, I briefly summarize the contributions I have made in this dissertation. I structure my analysis of my contributions via my original research questions, and analyze how different studies contributed to these contributions. This section also contains a fair deal of reflection on some of the implications of these studies which have compelled me, as well as some comments about where I intend to take my research next. I begin with the contributions.

8.1 Primary Contributions

RQ1 (*Structure*): What methods can we use to characterize the members, boundaries, and interrelations of different influencer communities? (Chapters 4, 5, 6, 7)

8.1.1 RQ1a: Tools for Measuring the Structure of Influencer Communities

Each of my studies was designed to illuminate different aspects of the structure of influencer communities in U.S. political discourse. Study 1 laid out a basic method for understanding these communities: the coengagement network. The coengagement network posited that influencers were drawn into relation with one another if they shared a similar audience, and thus similar conditions under which to generate *attention.* Community here was not typically defined as actual, personal relationships between influencers, but as an imagined community of shared circumstance: people

laboring under similar conditions, though perhaps without unity or even direct knowledge of each other. The coengagement network provided an empirical and visual tool for understanding the number of, membership within, and interrelation between different communities of influencers, perhaps most directly answering the research question laid out in **RQ1.** Study 1 showed how a detailed analysis of coengagement networks could reveal unique phenomena within influencer communities, including manufactured influence cartels that we termed followback networks; tenuous points of connection with influencer communities outside of mainstream U.S. political discourse, such as those associated with pop music or non-English communities; and those unique accounts who use posting strategies designed to maximize attention from multiple communities at once.

Study 1 proposed a method for measuring influencer communities at static timepoints, and so Study 2 proposed a framework for understanding how this tool can be used to study influencer communities longitudinally at two separate timepoints. As I remark in this study, the concept of "change" in influencer communities will depend on which aspects of these communities one wants to measure. I present a relatively simple change framework consisting of changes of membership of audiences, influencers themselves, changes in their activity levels, and changes in platform affordances which connect audiences to influencers. Changes in these qualities signal bulk changes in the conditions under which influencers must operate. As audiences grow or shrink, their demands and ability to give their attention change. I demonstrate that coengagement networks taken at different points but under similar circumstances (such as repeated elections within the same country) can facilitate effective measurement of these and other qualities — and provide useful correctives to public assumptions about the growth or decline of influencer communities online.

Study 3 elaborated on the structure of influencer communities by adding a valence to them. Communities could be modeled to have different types of relationships with each other, specifically negative relationships of antagonism and non-negative relationships of tolerance or approval.

This added dimension helped reveal, on a structural level, how some influencer communities were analogous to dominant publics, counterpublics (Fraser, 1990; Squires, 2002), and defensive publics (Jackson and Kreiss, 2023), due to the ways they used or did not use negativity to relate to other publics. The analysis of negativity also reveals relationships of dependence between influencer communities that mediated negativity based on ideological opposition. Influencer communities which control access to industrialized news, or rather the production of trusted information on a variety of topics, are a source of discourse upon which other influencer communities must depend for maintaining attention. In this light, the dynamics of how these dominant news-generating communities relate to other communities will always be distinct; to control the generation of information is a valuable resource in influencer economies.

Study 4 was a divergence from the first three studies, in that it defined a parallel structure to the influencer community assumed to be its product. This is the structure of consensus, which influencers collaboratively produce by selectively endorsing and rejecting pieces of evidence. Study 4 reprises the theme of negativity from Study 3, showing how the construction of knowledge in influencer communities is sometimes inseparable from the tearing down of knowledge from other communities. Influencers not only tell us what we should believe, but what we should not believe, i.e. what opposed influencers tell us. Seeing knowledge construction as a process that includes both selection and rejection widens the scope of any discourse community to necessarily include its opponents, to rise the level the analysis from the influencer *community* to, perhaps, the influencer *ecosystem*. This study begins a line of research into the connections between the daily world of influence and the discourses that their work creates, which I hope to continue in my post-dissertation work.

Public, Audience, or Community?

Throughout this dissertation, I have chosen the word *community* to refer to groups of influencers. My inspiration here comes from imagined communities and communities of practice, where would-be influencers come to a sense of themselves as influencers within a particular audience through the reading of metrics generated from those audiences. However, several other terms may be used to describe the activities I chronicle in these studies, which I briefly review here.

The first is concept of influencers as a public, or what Abidin (2021) has called the *refracted public*. Publics have typically been described in normative terms as spaces where private citizens come together to discuss political priorities and ultimately filter up policy imperatives to those that govern them. They are juxtaposed against mass media consumers, who consume information largely irrelevant to governing and whom are controlled by the interests of industrial and/or state actors. As is typical in the history of study on publics, these case studies reveal structures which operate somewhere in between. As Hund (2023, p. 183) has also noted in her own reflections on political influencers and their implications for democracy, it is likely that most of the influencers I examine are influencing out a genuine desire to inform their communities and contribute to democratic decision-making. Likewise, it is certain that many of the influencers I examine are motivated by economic, political, or affective goals which are not *immediately* relevant to governing. These motivations include the need to sustain an independent media career, or the need to win an election, or the need to retain the admiration of one's followers and avoid harassment from one's detractors. It is an intentional choice in my method to represent these influencers with widely different motivations as essentially similar nodes in a *community,* as regardless as what their aims are, they are all trying to work the same system to achieve them.

In the tradition of human-computer interaction literature, it could be said that the influencer economy — the technological innovations of social media platforms, the industrialization innovations in influencer marketing and metricization, the cultural innovations in the U.S. to understand influencers as a legitimate source of information — provides *affordances* for the normative operation

of a public. However, it also provides affordances for any number of other self-interested uses. In this sense, it is little different from the innovative media cultures that preceded it, such as the blogosophere, broadcast television, radio, or daily newspapers. While influencer communities have the potential to be publics in the Habermasian ideal, they need not be — as the self-interested COVID-19 influencers in Study 4 (Chapter 7) illustrate.

Another metaphor in the literature for these communities is that of an audience. Contemporary influencers occupy an indeterminate space between audiences and media. In some sense, they are audience members in that they frequently process messages from an (increasingly fractured) mass media apparatus. However, in another sense, as recent reporting on leading sources of news among U.S. youth shows (Robb, 2020), they increasingly *are* the mass media for a growing segment of the population. Early work on opinion leaders, those supposed intermediaries between the mass media and the public, often suggested that opinion leaders were merely particularly informed and popular audience members that shared their opinions with peers. By contrast, my studies' empirical results consistently reveal that influencers are some of the most popular sources of online engagement on Twitter, which includes many sources of mass and "mainstream" news. These studies portray, in other words, ongoing horizontalization of media and audiences, where the generation of trusted information is placed on an equal footing with the subsequent interpretation and re-distribution of that information. As other scholars of networked media have implied, the linear model of communication from media to audiences, is seriously threatened, to the extent it ever fully existed (Billard and Moran, 2020). My intent is for the coengagement graph method to stress this horizontalization, to put mainstream news sources on equal footing with platform-built influencers in the contemporary *networked* model of media communication.

We can similarly contrast my use of the word "community" with research on fandom communities, which intertwine heavily with scholarship on audiences. Fandom communities can have an eerie resonance with the influencer communities I portray in my studies. In political discourse on Twitter, politicians and other well-known political figures are speculated upon as if they are

characters in a dramatic soap opera — not too different from the way that Usenet users discussed literal soap operas at the advent of the internet (Baym, 1999). As with opinion leaders, however, there is an increasing ontological confusion between which users are fans and which users are the objects of their attention. Study 3 (Chapter 6) shows that, especially in pro-Trump communities, influencers themselves become the subject of drama and negative speculation, alongside the politicians of which they could be said to be fans. If influencer communities are scaled-up fan communities, they are fan communities which have the ability to generate their own momentum: to generate new media celebrities from the process of speculating on old media celebrities. The empirical finding that these networks are typically centered around platform-built influencer accounts, rather than those for politicians or celebrity journalists, supports the notion that fans are becoming their own icons.

Ultimately, I settle on the metaphor of the imagined community (Anderson, 1983). As happens with all media innovations, the past few decades have been characterized by a conspicuous anxiety about how to manage information and visibility in a media environment that creates and distributes information at an unprecedented pace. There is a deep vein of research, for example, on how users studiously observe and discuss contemporary social media platforms in order to understand how they work, and particularly by which methods they may see and be *seen* (DeVito, Gergle, and Birnholtz, 2017; Bishop, 2019; Simpson, Hamann, and Semaan, 2022). The imagined audience, or rather the metrics, signs, and gossip which reflect the imagined audience, are a sort of shared text which generates a shared identity in the contemporary influencer media ecosystem. As these studies' empirical findings consistently show, different target audiences generate metrics in different ways and in response to different things, structuring the cultures of influencers who scrutinize them in unique ways. While I do not intend to displace other metaphors for the group activities of influencers, understanding the structure of groups of influencers as imaginaries can help guide our focus towards how influencers are perceiving their own worlds — and how their perceptions create effects reaching far outside their own quest for attention.

8.1.2 RQ1b: The Empirical Structure of U.S. Political Discourse on Twitter

Study 1 also laid out the basic schematic of which communities, empirically, were dominant in U.S. political discourse as defined through the lens of the 2020 United States election. Twitter holds a tripartite structure in this context, with an approximately pro-Trump, mainstream progressive, and radical socialist community of influencers. The act of naming these communities is fraught, and is naturally powerful due to my status as a researcher in an R1 institution. Readers will observe that I have iterated through different naming conventions between studies, as my understanding of these communities have changed over time. In this final subsection, I provide reflections on the process of understanding these different communities, and the roles they play in U.S. Twitter discourse.

Before I summarize, I attend to the difficulties of naming. Influencers who share a community, per my definition, do not necessarily themselves share values or require any self-determined associative identity, making them difficult to name directly. Many communities will include influencers with opposing values, but who can appeal to different desires of their shared audiences at different times. People also hold strong expectations of and normative judgments of political communities, making it difficult to produce a name that will not invoke prior associations. Names also sensitize readers to which features of a group the authors believe are most salient, which could be a group's values, practices, or associations with other groups.

Mainstream Progressive Influencers

What I have titled the mainstream progressive cluster in Study 3, and formerly the pro-Biden and left-leaning cluster in Studies 1 and 2, embodies many of these difficulties. While clearly focused around U.S. president Joseph Biden, the community admits much criticism of his policies and leadership at any given moment, indicating an underlying value structure which Biden may or may not align with at any given moment. And yet, because he is the presumptive candidate of the U.S.'s leading progressive party, support for him is often a binding principle between members of

this community. The broader term "progressive" appears to fit this community's discourse better than liberal or Democratic (party), given its conditional support of policies like the defunding of the policing system, which is a liberal/Democratic institution.

Perhaps most controversial is the label "mainstream," given the positive associations with the word in popular discourse. Mainstream does not necessarily signify representing a majority of people within a public, though it appears that this community has the highest membership and more often correlates with the mostly widely-held values and policy beliefs among U.S. residents. Rather, it signifies an association with a dominant, central public which among other things holds a privileged place in the generation and distribution of widely-trusted knowledge. The mainstream cluster contains the most news organizations, often the most well-funded and widely-read news organizations, and news organizations which more typically (but not always) follow journalistic norms of objectivity and sourcing. It also contains governmental and scientific institutions and actors, who are themselves bearers and generators of a certain kind of valorized truth.

The mainstream defines itself in contrast to the "alternative," a phrase recently popularized by the rise of the authoritarian "alt-right," but with a much longer history of counter-news and truth-making institutions which serve as attempt correctives to the "mainstream" perspective (Lewis, 2018; Starbird et al., 2018; Holt, Ustad Figenschou, and Frischlich, 2019). Though mainstream knowledge sources are perhaps the most widely trusted sources, especially among the powerful, they still represent a typically liberal ideology which at many points appears to misrepresent the experiences of marginalized groups and entrench the interests of the powerful (Herman and Chomsky, 1988; Mejia, Beckermann, and Sullivan, 2018). Indeed, "alternative" sources of truth can be more consonant with marginalized experiences and positive change (Michaeli, 2016), though of course they can also be sources of fascistic and authoritarian information which entrench powerful groups even further (Sandoval and Fuchs, 2010; Beers, Wilson, and Starbird, 2022). Thus, while many people associate the moniker "mainstream" with a positive valence in our recent age of information disorder (Aftergut and Lacovara, 2023), I use it with a more ambiguous valence in

the mainstream progressive cluster.

Pro-Trump Influencers

The second-largest influencer community in these case studies is what I term the pro-Trump influencer community. Generally associated with right-leaning politics, one could argue that ideological labels such as right-leaning, conservative, far-right, fascist, white nationalist, misogynist, authoritarian, and populist. I agree that all of these labels can apply to individual members of this community at the very least, and many to the community as a whole. I have, for the moment, settled on the pro-Trump label due to the community's dependence on former president Donald Trump for the orientation of all its discourse. Study 1 demonstrated that Trump, when active on the platform, holds a disproportionate preeminence in this community completely unmatched by the most popular users in other communities. Study 3 demonstrated that resistance to Trump as a person, rather than any particular value or policy, was one of the driving factors inducing norm correction from one's co-influencers. While the focus on a single person may itself represent a sort of authoritarian or totalitarian project, it seems more descriptive to name the community's complete dependence on the actions and character of a single person.

Radical Socialist Influencers

The third, much smaller, influencer community is what I have titled as of Study 3 as radical socialist influencers. While "radical" and "socialist" are often seen as negative terms in popular U.S. discourse, and especially conservative U.S. discourse, I see them as neither.

I take radical here to mean a commitment to pursuing political solutions outside the context of existing institutions. I use the term not as an absolute set of policy positions, but rather a posture towards one's existing context, in this case the U.S. context in the early 2020s (Pugh, 2009). People with radical beliefs in the U.S. may hold non-radical beliefs in other countries, and beliefs which were previously radical may not be integrated into mainstream institutions. And while

radical may sometimes signify a politics I personally sympathize with, it can take form in ideologies which are often in direct conflict — see the popular term "radical right" for a comparison in this respect (Rydgren, 2007; Pirro, 2023). Similarly, radical politics need not signify a lack of institutional power, marginalized subject position, or anti-capitalism politics (see Sandoval and Fuchs, 2010; Jackson and Kreiss, 2023 for discussions of counter-institutional discourse with more normative frames).

Similarly, I understand the phrase "socialist" here to not imply any particular policy position or organizational affiliation. Many prominent members of this influencer community self-identify as some form of socialist, although some do not. Their labeling as socialist in this dissertation accordingly reflects their association with a certain popular discourse around the idea of socialism, rather than a historical connection with socialist movements or ideas.

Their presence in this data is normatively ambiguous, and challenges many frameworks for understanding counter-mainstream publics. For example, they generally advocate for many progressive radical issues: the abolition of policing; the return of land from settler to Native Americans; the free provision of healthcare, housing, and other basic survival needs; the dismantling of the U.S. military infrastructure as currently constituted, if not entirely. However, they also take stances challenging progressive orthodoxy and supporting certain existing institutions, such as a support for a Russian government which many consider to be authoritarian (Beers, Wilson, and Starbird, 2022). While they in many cases advocate for redressing harms to marginalized populations, their outlook and composition is frequently based upon a rejection of identity-based organizing. It is thus difficult to place them strictly into the framework of counterpublics, defensive publics (Jackson and Kreiss, 2023), marginalized communities, or other terms frequently used to describe counter-mainstream media aggregates.

Their presence in this dissertation is a much-needed complicating factor in the understanding of U.S. political influencer communities. They are counter-mainstream, yet far more marginal on- and offline than pro-Trump influencers by most metrics. They are decidedly non-binary in their

policy positions, supporting mainstream progressive influencers in some cases, and pro-Trump influencers in others. In my future work, I hope to better understand their growth and impact on political discourse writ large.

From Belief Spectrums to Ideological Communities

I conclude with a reflection on representing structures of ideology as downstream of coengagement networks rather than, for example, continuous metrics of left-right ideology favored in most social media studies (Barberá, 2015; Benkler, Faris, and Roberts, 2018). I believe measuring ideology on social media platforms as downstream of influencer communities is advantageous for three reasons. The first is that ideology is non-binary and context-dependent. The radical socialist community is the greatest example of this, mixing policy beliefs from the other two communities in a way that defies a straightforward placement on a two-dimensional ideological spectrum. The second is that personal beliefs and ideologies are not stable, but are shaped by media, and in this case influencer media. Describing ideologies in terms of the media environments they filter down from can help alleviate stressful ideological inconsistencies that may arise due to unique dynamics of those media environments. The third is that influencer communities are composed of people, each of whom consists of a rich case study of the actualization of ideology in ever-shifting contexts, and reveals a certain richness to belief formation absent in point-space models of ideology. Ideology is manufactured by this ecosystem of influencers in a complex and unpredictable way, and I believe we should analyze ideology in a way that can provide better access to these complexities.

8.1.3 RQ2: Strategies Within Influencer Communities

RQ2 (*Strategy*): How can we characterize the strategies that influencers use in the aggregate to attract sustained attention, with respect to their own communities and their peer communities? (Chapters 4, 6, 7)

My second contribution revolves around the strategies that different influencers use to maintain attention and audience share. Study 1, in its exploration of the unique features of coengagement networks, illustrated many of these strategies. Followback networks create artificial retweet cartels to increase their public profiles and hopefully attract organic attention; non-English satellite influencers translate tweets from dominant publics to maintain audience share in their own communities; "non-partisan" news sources selectively craft headlines and tweets to appeal to different audiences without entirely alienating one or the other. Study 1 portrays a vibrant galaxy of strategic influencing niches, uniquely enmeshed in the technological affordances and cultural context of U.S. online political discussions.

Study 3 furthered an analysis of influencer strategy by introducing the use of negativity. It clarified that being the subject of negative attention, regardless of popular narratives around sensationalism and the oxygen of amplification (Phillips, 2018), is not necessarily helpful for building a following. However, it does seem to be used for norm enforcement both within and without communities, and that norm enforcement aimed at out-groups is a particularly effective strategy for building a following. The results of this study suggest that the maintenance of attention online is essentially a symbiotic enterprise with one's opposing influencers; each gains attentional capital through the denigration of those outside their community. It also suggests that violating the norms of one's community can be perilous; both an opportunity for your co-influencers to gain attention, and potentially an opportunity for you to lose yours.

This potential symbiosis is further explored in Study 4, which studies how influencers gather (scientific) evidence to present and presumably curry favor from their audiences. Again, we see that influencers, especially those posed against the mainstream, consistently use negative citation of the opposing sides' materials to bolster their own arguments. The relationship here may be less symbiotic than defensive, guarding against opposing influencers' attempts to reduce their legitimacy with "better" evidence. But the point remains the same: influencer strategies are interactional with other influencers, and to prosper in an influencer economy often involves not only

following the norms of one's influencer community, but consciously denigrating the members of out-communities.

The Influencer Theater

In my earlier comments on the methodology of coengagement networks (Chapter 3), I referred to the practice of representing influencers through the medium of shared audiences as essentially representing, per Goffman (1959), a *performance.* While influencer communication is often presented as a singular performance, Studies 3 and 4 in particular show that their performances are frequently interactional. In presenting this work, I have often used the metaphor of the *influencer theater.* Influencers sometimes monologue, sometimes engage in crowdwork, sometimes perform duets and group scenes (Lewis, 2018), and alternately play the hero and villain in dramatized scenes of conflict for their audiences (Christin and Lewis, 2021; Lewis and Christin, 2022). I do not doubt that many influencers from between politically-opposed communities actually do dislike each other and resent each others' values, but it is also clear that certain influencers — particularly those in relatively dominant subject positions — benefit in the attention economy from the performance of conflict with their dramatized nemeses.

The fact that influencers' benefit from their feuds does not mean these feuds are frivolous or self-interested. The history of activism is partly the history of publicized conflict, and there is much scholarship on the productive uses of conflict in supporting marginalized groups (Lorde, 1997; Nakamura, 2015; Clark, 2020; Peterson-Salahuddin, 2022). However, others, such as (Duffy, Miltner, and Wahlstedt, 2022), have noted that negativity and conflict even in the name of ostensibly progressive causes can be ultimately regressive and reproductive of harmful power dynamics. Understanding when collaboration — and conflict — is activism and norm enforcement, and when it is simply more performances in the influencer theater, is a priority for my future research on the strategies of influencers in influencer communities.

8.1.4 RQ3: The Products of Influence

RQ3 (*Effects*): How can we measure what influencer communities do, and particularly how their activities structure online political discourse? (Chapters 6, 7)

The third contribution centers around how influencers shape online political discourse. While Studies 1 and 2 of this dissertation focus primarily on structure and change in influencer communities, Studies 3 and 4 focused more directly on the content of influencers' utterances, and what effect they may be having on broader political discourse.

In Study 3, we see how political influencers use negative commentary to enforce ideological norms within their own communities. Pro-Trump communities chastise those who lack in their support for Trump and his policies, while mainstream progressive communities chastise those who, for example, oppose certain COVID-19 mitigation measures. It is clear then that influencer communities perform a sort of self-moderation, attempting to reduce the exposure of viewpoints which deviate from popular consensus within their group.

Are influencers then ideological umpires, keeping their fellow influencers and audiences from straying too far from community dogma? Or are they responding to incentives from their audiences, spotting a rich opportunity for positive feedback at their co-influencer's expense? This study, and dissertation, cannot yet say. Study 3 indicated that posts which negatively targeted co-influencers performed much worse than non-negative posts targeting co-influencers, suggesting no direct, immediate benefit to enforcing community norms. However, there may be indirect long-term benefits to enforcing community norms, specifically by shoring up one's appearance as an true, *authentic* believer in community dogma. Literature on influencers has long stressed the pressure that influencers feel to appear authentic, and how audiences' perception of authenticity affect their feelings towards influencers (Duffy, 2017; Abidin, 2018; Pöyry et al., 2021; Duffy, Miltner, and Wahlstedt, 2022; Hund, 2023; Arnesson, 2023). I intend to pursue research understanding incentives for influencers to demonstrate authenticity as a mechanism of ideological control in

heavily-politicized influencer communities.

While Study 3 analyzed the shaping of political discourse only indirectly, Study 4 concerned itself directly with how influencers shape our perceptions of scientific evidence, and consequently of which public health interventions their audiences should endorse or criticize. We suggested in Study 4 that two of the principle actions of influencers are to aggregate information into an evidence-base, or *consensus*, on the one hand, and to specifically denigrate information outside of that evidence base as untrustworthy. We used a method similar to the coengagement graph adjusted for negative ties to show consensus formation in how influencers talked about whether masks prevented COVID-19, and consequently whether mask-wearing should be policy. Medical influencer communities effectively cleaved the information world in two, performing boundary-work (Gieryn, 1983) to simplify an overwhelming situation for their audiences. It is this process of aggregation, denigration, and resulting binarization which I suggest is one of the essential mechanisms by which influencers shape their audiences perceptions' of political dilemmas.

Study 4, in its construction, was agnostic to the popular legitimacy of the influencers it studied. A science journalist employed by the renowned University of California San Francisco health system was considered with the same methods as the career alternative media conspiracy theorist. Its results suggest, however, that the legitimacy of an influencer's status within a media ecosystem may affect the extent to which they use aggregation and denigration. As we discuss in Study 4, alternative media influencers who oppose the mainstream scientific consensus on masks appeared to explicitly denigrate that consensus more than vice versa. This suggests different strategies, which naturally lead to different results in political discourse. Alternative media and counterpublics may need to spend more of their time using negativity to reframe mainstream belief as untrustworthy, while mainstream and dominant publics may simply ignore these alternative publics due to their marginality. I suggest (and plan to pursue) additional research into the interaction between the shaping of political discourse by influencers and the marginality of the influencer communities in which they reside.

8.2 Future Directions

One of this dissertation's projects was to create a new and hopefully productive definition of the influencer and influencer community for future research. Aspects of that definition developed over time, and my understanding of influencers' motivations was one of the last pieces to click, which led it to be neglected in these studies. Much has been written and theorized about the way that influencers generate income from their art, particularly in media studies. As I continue my research, I plan to continue used mixed-methods to understand how influencers engaging in self-consciously political discourse, such as the influencers I study in this discourse, generate income, and how their strategies affect mainstream political discourse in online and offline environments.

Research into profiting from explicitly politicized influence is blossoming in the past five years. The COVID-19 pandemic exposed the naked profiteering of wellness and anti-vaccine conspiracy influencers, and generated efforts to document their particular monetization strategies (Herasimenka et al., 2022; Broniatowski et al., 2023; Moran, Swan, and Agajanian, 2024). While many have noted the profit motive in influencers' political rhetoric, particularly among far-right influencers, descriptions of the systems by which these influencers transform political rhetoric into income are nascent (Gregorio and Goanta, 2022; Munger and Phillips, 2022). Research into how economic incentive structures in the influencer paradigm systematically shape political discourse are rarer still (though see Hund, 2023, "The Industry Becomes Boundaryless"), even as plenty such similar studies exist for other media paradigms (Lippmann, 1922; Adorno and Horkheimer, 1944; Herman and Chomsky, 1988; Weigel, 2022).

Accordingly, my next step is to begin a process of mapping political discourse generated by influencers to metrics of valuation increasingly produced by rising media platforms. Rising media companies like Spotify, YouTube, and Substack increasingly generate data, both quantitative and qualitative, on which discourses are most profitable within an influencer-driven digital media ecosystem. A growing body of research shows that influencers themselves pay assiduous attention to these platform policies and metrics in their participation in the culture industry (Nieborg,

Duffy, and Poell, 2020). We, too, can learn from these platform environments, and in the process develop an understanding about how the profit-seeking practices and economic constraints placed on influencers delimit the bounds of political discourse writ large.

8.3 Conclusion

This dissertation conducts a study of influencer communities, which I define as large groups of influencers subject to similar audiences on the same platform. It is partly methodological and descriptive, in that it proposes several methods for defining and describing these communities, and partly empirical and explanatory, in that it describes an extensive case study of influencer communities in U.S. political discourse on Twitter. I contribute to our understanding of influencers and their communities in at least three ways. First, I propose methods from network science for analyzing and visualizing the structure of, long-term change in, and patterns of association between influencer communities, and provide empirical results in the U.S. context. Second, I use the case study perspective to describe how influencers use certain strategies, such as retweet cartels or social shaming, to attract attention and build their followings. Third, I describe some of the effects that influencers have on online political discourse writ large, particularly in the enforcement of ideological norms and the gathering of evidence for those norms. I have concluded this dissertation by repeatedly discussing power, and how akin to research on counterpublics and defensive publics, power gradients between different communities affects the strategies they must pursue and the potential effects they can have on public discourse.

Bibliography

Abdalla, Mohamed et al. (2023). "The Elephant in the Room: Analyzing the Presence of Big Tech in Natural Language Processing Research". In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13141–13160. DOI: 10.18653/v1/2023.acl-long.734. URL: http://arxiv.org/abs/2305.02797 (visited on 05/04/2024).

Abidin, Crystal (Apr. 2016). ""Aren't These Just Young, Rich Women Doing Vain Things Online?": Influencer Selfies as Subversive Frivolity". en. In: *Social Media + Society* 2.2, p. 2056305116641342. ISSN: 2056-3051. DOI: 10.1177/2056305116641342. URL: https://doi.org/10.1177/2056305116641342 (visited on 04/20/2024).

– (Apr. 2017). "#familygoals: Family Influencers, Calibrated Amateurism, and Justifying Young Digital Labor". en. In: *Social Media + Society* 3.2, p. 2056305117707191. ISSN: 2056-3051. DOI: 10.1177/2056305117707191. URL: https://doi.org/10.1177/2056305117707191 (visited on 05/07/2024).

– (2018). *Internet celebrity: Understanding fame online*. Emerald Group Publishing.

– (2019). "Victim, Rival, Bully: Influencers' Narrative Cultures Around Cyberbullying". en. In: *Narratives in Research and Interventions on Cyberbullying among Young People*. Ed. by Heidi Vandebosch and Lelia Green. Cham: Springer International Publishing, pp. 199–212. ISBN: 978-3-030-04960-7. DOI: 10.1007/978-3-030-04960-7_13. URL: https://doi.org/10.1007/978-3-030-04960-7_13 (visited on 03/24/2023).

– (Jan. 2021). "From "Networked Publics" to "Refracted Publics": A Companion Framework for Researching "Below the Radar" Studies". en. In: *Social Media + Society* 7.1, p. 2056305120984458. ISSN: 2056-3051. DOI: 10.1177/2056305120984458. URL: https://doi.org/10.1177/2056305120984458 (visited on 06/04/2024).

Abilov, Anton et al. (Apr. 2021). "VoterFraud2020: a Multi-modal Dataset of Election Fraud Claims on Twitter". In: *arXiv:2101.08210 [cs]*. (Visited on 09/07/2021).

Abokhodair, Norah et al. (Apr. 2024). "Opaque algorithms, transparent biases: Automated content moderation during the Sheikh Jarrah Crisis". en. In: *First Monday*. ISSN: 1396-0466. DOI: 10.5210/fm.v29i4.13620. URL: https://firstmonday.org/ojs/index.php/fm/article/view/13620 (visited on 05/04/2024).

Acker, Amelia (Nov. 2018). *Data craft: the manipulation of social media metadata*. en. Report. Data & Society Research Institute. URL: https://apo.org.au/node/202091 (visited on 05/08/2023).

Adorno, Theodor W. and Max Horkheimer (1944). "The Culture Industry: Enlightenment as Mass Deception". In: *Dialectic of enlightenment*. Vol. 15. Verso. URL: https://books.google.com/books?hl=en\&lr=\&id=lwVjsKcHW7cC\&oi=fnd\&pg=PR9\&dq=dialectic+of+enlightenment\&ots=8UCo05KRss\&sig=AzVGi2A6g5c_QVLk95PRPRVKjSw (visited on 05/05/2024).

Aftergut, Dennis and Phillip Allen Lacovara (Aug. 2023). "The media is winning the war against "fake news"". en. In: *Salon*. URL: https://www.salon.com/2023/08/23/the-mainstream-media-is-winning-the-against-fake-news/ (visited on 05/02/2024).

Alatas, Vivi et al. (Feb. 2019). *When Celebrities Speak: A Nationwide Twitter Experiment Promoting Vaccination In Indonesia*. Working Paper. DOI: 10.3386/w25589. URL: https://www.nber.org/papers/w25589 (visited on 12/16/2022).

AlDayel, Abeer and Walid Magdy (July 2021). "Stance detection on social media: State of the art and trends". en. In: *Information Processing & Management* 58.4, p. 102597. ISSN: 0306-4573. DOI: 10.1016/j.ipm.2021.102597. URL: https://www.sciencedirect.com/science/article/pii/S0306457321000960 (visited on 04/03/2023).

Allgaier, Joachim (2019). "Science and Environmental Communication on YouTube: Strategically Distorted Communications in Online Videos on Climate Change and Climate Engineering". In: *Frontiers in Communication* 4. ISSN: 2297-900X. URL: https://www.frontiersin.org/articles/10.3389/fcomm.2019.00036 (visited on 08/26/2022).

Analytics, Clarivate (2017). "Web of science core collection". In: *Citation database. Web of Science*.

Andersen, Kim, Adam Shehata, and Dennis Andersson (May 2023). "Alternative News Orientation and Trust in Mainstream Media: A Longitudinal Audience Perspective". In: *Digital Journalism* 11.5, pp. 833–852. ISSN: 2167-0811. DOI: 10.1080/21670811.2021.1986412. URL: https://doi.org/10.1080/21670811.2021.1986412 (visited on 05/02/2024).

Anderson, Benedict Richard O'Gorman (1983). *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. en. Verso. ISBN: 978-0-86091-546-1.

Ang, Ien (1991). *Desperately Seeking the Audience*. Oxford, UNITED KINGDOM: Taylor & Francis Group. ISBN: 978-0-203-13334-7. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=165716 (visited on 06/03/2024).

Angwin, Julia (May 2023). "Few Are Addressing One of Social Media's Greatest Perils". en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2023/05/06/opinion/fear-speech-social-media.html (visited on 05/01/2024).

Arif, Ahmer (2020). "Troubling Matters: Examining the Spread of Misinformation and Disinformation on Social Media during Mass Disruption Events". English. Ph.D. United States – Washington: University of Washington. URL: https://www.proquest.com/docview/2492246343/abstract/D1E57B555F58406CPQ/1 (visited on 05/08/2024).

Arif, Ahmer, Leo Graiden Stewart, and Kate Starbird (Nov. 2018). "Acting the Part: Examining Information Operations Within #BlackLivesMatter Discourse". en. In: *Proceedings of the ACM on Human-Computer Interaction* 2.CSCW, pp. 1–27. DOI: 10.1145/3274289. (Visited on 05/06/2020).

Arnesson, Johanna (Apr. 2023). "Influencers as ideological intermediaries: promotional politics and authenticity labour in influencer collaborations". en. In: *Media, Culture & Society* 45.3, pp. 528–544. ISSN: 0163-4437. DOI: 10.1177/01634437221117505. URL: https://doi.org/10.1177/01634437221117505 (visited on 04/20/2024).

Arriagada, Arturo and Sophie Bishop (Dec. 2021). "Between Commerciality and Authenticity: The Imaginary of Social Media Influencers in the Platform Economy". In: *Communication, Culture and Critique* 14.4, pp. 568–586. ISSN: 1753-9129. DOI: 10.1093/ccc/tcab050. URL: https://doi.org/10.1093/ccc/tcab050 (visited on 05/07/2024).

Arsenyan, Jbid and Agata Mirowska (Nov. 2021). "Almost human? A comparative case study on the social media presence of virtual influencers". In: *International Journal of Human-Computer Studies* 155, p. 102694. ISSN: 1071-5819. DOI: 10.1016/j.ijhcs.2021.102694. URL: https://www.sciencedirect.com/science/article/pii/S1071581921001129 (visited on 05/07/2024).

Asher, Steven R. and Kenneth A. Dodge (July 1986). "Identifying children who are rejected by their peers: Developmental Psychology". In: *Developmental Psychology* 22.4, pp. 444–449. ISSN: 0012-1649. DOI: 10.1037/0012-1649.22.4.444. URL: http://offcampus.lib.washington.edu/login?url=https://search.ebscohost.com/login.aspx?direct=true\&db=pdh\&AN=1986-28929-001\&site=ehost-live (visited on 04/13/2024).

Badawy, Adam, Emilio Ferrara, and Kristina Lerman (Aug. 2018). "Analyzing the Digital Traces of Political Manipulation: The 2016 Russian Interference Twitter Campaign". In: *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 258–265. DOI: 10.1109/ASONAM.2018.8508646.

Bai, Nina (June 2020). *Still Confused About Masks? Here's the Science Behind How Face Masks Prevent Coronavirus*. en. URL: https://www.ucsf.edu/news/2020/06/417906/still-confused-about-masks-heres-science-behind-how-face-masks-prevent (visited on 11/02/2022).

Bailey, Moya (May 2021). *Misogynoir Transformed: Black Women's Digital Resistance*. New York University Press. ISBN: 978-1-4798-0339-2. DOI: 10.18574/nyu/9781479803392.001.0001. URL: https://www.degruyter.com/document/doi/10.18574/nyu/9781479803392.001.0001/html (visited on 04/21/2024).

Baker, Stephanie Alice (Feb. 2022). "Alt. Health Influencers: how wellness culture and web culture have been weaponised to promote conspiracy theories and far-right extremism during the COVID-19 pandemic". en. In: *European Journal of Cultural Studies* 25.1, pp. 3–24. ISSN: 1367-5494. DOI: 10.1177/13675494211062623. URL: https://doi.org/10.1177/13675494211062623 (visited on 04/20/2024).

Bakshy, Eytan et al. (Feb. 2011). "Everyone's an influencer: quantifying influence on twitter". In: *Proceedings of the fourth ACM international conference on Web search and data mining*. WSDM '11. New York, NY, USA: Association for Computing Machinery, pp. 65–74. ISBN: 978-1-4503-0493-1. DOI: 10.1145/1935826.1935845. URL: https://doi.org/10.1145/1935826.1935845 (visited on 12/08/2022).

Balnaves, Mark, Tom O'Regan, and Ben Goldsmith (2011). *Rating the Audience : The Business of Media*. English. Bloomsbury Academic. ISBN: 978-1-84966-461-5 978-1-84966-460-8. DOI: 10.5040/9781849664622. URL: https://directory.doabooks.org/handle/20.500.12854/92860 (visited on 06/03/2024).

Banet-Weiser, Sarah (2012). *Authentic(tm): The Politics of Ambivalence in a Brand Culture*. New York, UNITED STATES: New York University Press. ISBN: 978-0-8147-3937-2. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=865546 (visited on 05/07/2024).

Barabási, Albert-László (Mar. 2013). "Network science". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371.1987, p. 20120375. DOI: 10.1098/rsta.2012.0375. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2012.0375 (visited on 04/23/2024).

Barbala, Astri Moksnes (Aug. 2023). "Reassembling #MeToo: Tracing the techno-affective agency of the feminist Instagram influencer". en. In: *Convergence*, p. 13548565231191261. ISSN: 1354-8565. DOI: 10.1177/13548565231191261. URL: https://doi.org/10.1177/13548565231191261 (visited on 04/21/2024).

Barberá, Pablo (2015). "Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data". In: *Political analysis* 23.1, pp. 76–91. URL: https://www.cambridge.org/core/journals/political-analysis/article/birds-of-the-same-feather-tweet-together-bayesian-ideal-point-estimation-using-twitter-data/91E37205F69AEA32EF27F12563DC2A0A (visited on 11/06/2023).

Barberá, Pablo et al. (Nov. 2019). "Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data". en. In: *American Political Science Review* 113.4, pp. 883–901. ISSN: 0003-0554, 1537-5943. DOI: 10.1017/S0003055419000352. URL: https://www.cambridge.org/core/product/identifier/S0003055419000352/type/journal_article (visited on 10/20/2023).

Basol, Melisa et al. (Jan. 2021). "Towards psychological herd immunity: Cross-cultural evidence for two prebunking interventions against COVID-19 misinformation". en. In: *Big Data & Society* 8.1. ISSN: 2053-9517. DOI: 10.1177/20539517211013868. URL: https://doi.org/10.1177/20539517211013868 (visited on 08/29/2022).

Bastian, Hilda (Jan. 2023). *Quote Tweeting: Over 30 Studies Dispel Some Myths*. en-US. URL: https://absolutelymaybe.plos.org/2023/01/12/quote-tweeting-over-30-studies-dispel-some-myths/ (visited on 04/13/2024).

Bastian, Mathieu, Sebastien Heymann, and Mathieu Jacomy (Mar. 2009). "Gephi: an open source software for exploring and manipulating networks". en. In: *Third international AAAI conference on weblogs and social media*. Vol. 3. 1, pp. 361–362. URL: https://ojs.aaai.org/index.php/ICWSM/article/view/13937 (visited on 10/25/2022).

Baym, Nancy K. (1999). *Tune in, Log On: Soaps, Fandom, and Online Community*. Thousand Oaks, UNITED STATES: SAGE Publications, Incorporated. ISBN: 978-1-4522-2166-3. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=996454 (visited on 06/03/2024).

– (Sept. 2013). "Data not seen: The uses and shortcomings of social media metrics". en. In: *First Monday*. ISSN: 1396-0466. DOI: 10.5210/fm.v18i10.4873. URL: https://firstmonday.org/ojs/index.php/fm/article/view/4873 (visited on 06/03/2024).

Beers, Anna (June 2021). "Misinterpretation and Ambiguity in Public-Facing Network Visualizations: A Case Study". In: *Networks 2021: Networked Justice Satellite*.

– (June 2023). "Negative Influencers in Social Networks: From Targeted Harassment to Controversy Cultivation". In: *International Network for Social Network Analysis (Sunbelt)*.

Beers, Anna, Tom Wilson, and Kate Starbird (Sept. 2022). "The Demographics of an International Influence Operation Affecting Facebook Users in the United States". en. In: *Journal of Online Trust and Safety* 1.4. ISSN: 2770-3142. DOI: 10.54501/jots.v1i4.55. URL: https://tsjournal.org/index.php/jots/article/view/55 (visited on 12/08/2022).

Beers, Anna et al. (2021a). "Rejecting Science with Science: Boundary-Work in Anti-Mask Twitter Reply Threads During COVID-19". In: *AoIR Selected Papers of Internet Research*.

Beers, Anna et al. (2021b). "The Firestarting Troll, and Designing for Abusability". In: *International AAAI Conference on Web and Social Media, Information Credibility & Alternative Realities in Troubled Democracies Workshop*.

Beers, Anna et al. (June 2023a). "Followback Clusters, Satellite Audiences, and Bridge Nodes: Co-engagement Networks for the 2020 US Election". In: *Proceedings of the International AAAI Conference on Web and Social Media*.

Beers, Anna et al. (Sept. 2023b). "Selective and deceptive citation in the construction of dueling consensuses". In: *Science Advances* 9.38, eadh1933. DOI: 10.1126/sciadv.adh1933. URL: https://www.science.org/doi/10.1126/sciadv.adh1933 (visited on 10/16/2023).

Behrendt, Stefan (Mar. 2023). *lm.beta*. URL: https://cloud.r-project.org/web/packages/lm.beta/lm.beta.pdf (visited on 05/05/2024).

Benkler, Yochai, Robert Faris, and Hal Roberts (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. eng. New York: Oxford University Press. ISBN: 978-0-19-092362-4. DOI: 10.1093/oso/9780190923624.001.0001. URL: https://oxford.universitypressscholarship.com/10.1093/oso/9780190923624.001.0001/oso-9780190923624 (visited on 04/28/2022).

Benkler, Yochai et al. (Oct. 2015). "Social Mobilization and the Networked Public Sphere: Mapping the SOPA-PIPA Debate". In: *Political Communication* 32.4, pp. 594–624. ISSN: 1058-4609. DOI: 10.1080/10584609.2014.986349. URL: https://doi.org/10.1080/10584609.2014.986349 (visited on 05/16/2023).

Benkler, Yochai et al. (Oct. 2020). *Mail-In Voter Fraud: Anatomy of a Disinformation Campaign*. en. SSRN Scholarly Paper ID 3703701. Rochester, NY: Social Science Research Network. DOI: 10.2139/ssrn.3703701. (Visited on 01/16/2022).

Bennett, W. Lance and Shanto Iyengar (Dec. 2008). "A New Era of Minimal Effects? the Changing Foundations of Political Communication". In: *Journal of Communication* 58.4, pp. 707–731. ISSN: 0021-9916. DOI: 10.1111/j.1460-2466.2008.00410.x. URL: https://doi.org/10.1111/j.1460-2466.2008.00410.x (visited on 04/17/2024).

Bennett, W. Lance and Jarol B. Manheim (Nov. 2006). "The One-Step Flow of Communication". en. In: *The ANNALS of the American Academy of Political and Social Science* 608.1, pp. 213–232. ISSN: 0002-7162. DOI: 10.1177/0002716206292266. URL: https://doi.org/10.1177/0002716206292266 (visited on 12/12/2022).

Billard, Thomas J. (2023). *Voices for transgender equality: Making change in the networked public sphere*. Oxford University Press. URL: https://scholar.google.com/scholar?cluster=16103068048388176317\&hl=en\&oi=scholarr (visited on 04/17/2024).

Billard, Thomas J and Rachel E Moran (May 2020). "Networked political brands: consumption, community and political expression in contemporary brand culture". en. In: *Media, Culture & Society* 42.4, pp. 588–604. ISSN: 0163-4437. DOI: 10.1177/0163443719867301. URL: https://doi.org/10.1177/0163443719867301 (visited on 04/20/2024).

Bishop, Sophie (Nov. 2019). "Managing visibility on YouTube through algorithmic gossip". en. In: *New Media & Society* 21.11-12, pp. 2589–2606. ISSN: 1461-4448. DOI: 10.1177/1461444819854731. URL: https://doi.org/10.1177/1461444819854731 (visited on 06/04/2024).

– (Jan. 2021). "Influencer Management Tools: Algorithmic Cultures, Brand Safety, and Bias". en. In: *Social Media + Society* 7.1, p. 20563051211003066. ISSN: 2056-3051. DOI: 10.1177/20563051211003066. URL: https://doi.org/10.1177/20563051211003066 (visited on 06/04/2024).

Bishop, Sophie (Oct. 2023). "Influencer creep: How artists strategically navigate the platformisation of art worlds". en. In: *New Media & Society*, p. 14614448231206090. ISSN: 1461-4448. DOI: 10.1177/14614448231206090. URL: https://doi.org/10.1177/14614448231206090 (visited on 04/20/2024).

Block, Per (Jan. 2015). "Reciprocity, transitivity, and the mysterious three-cycle". en. In: *Social Networks* 40, pp. 163–173. ISSN: 0378-8733. DOI: 10.1016/j.socnet.2014.10.005. URL: https://www.sciencedirect.com/science/article/pii/S0378873314000586 (visited on 05/08/2023).

Blondel, Vincent D. et al. (2008). "Fast unfolding of communities in large networks". In: *Journal of statistical mechanics: theory and experiment* 2008.10, P10008.

Bond, Shannon (Aug. 2023). "Elon Musk sues disinformation researchers, claiming they are driving away advertisers". en. In: *NPR*. URL: https://www.npr.org/2023/08/01/1191318468/elon-musk-sues-disinformation-researchers-claiming-they-are-driving-away-adverti (visited on 04/16/2024).

Bonilla, Yarimar and Jonathan Rosa (2015). "#Ferguson: Digital protest, hashtag ethnography, and the racial politics of social media in the United States". en. In: *American Ethnologist* 42.1, pp. 4–17. ISSN: 1548-1425. DOI: 10.1111/amet.12112. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/amet.12112 (visited on 05/04/2024).

Booker, Brakkton (Jan. 2021). "Facebook Removes 'Stop The Steal' Content; Twitter Suspends QAnon Accounts". en. In: *NPR*. URL: https://www.npr.org/sections/insurrection-at-the-capitol/2021/01/12/956003580/facebook-removes-stop-the-steal-content-twitter-suspends-qanon-accounts (visited on 08/04/2023).

Boutron, Isabelle and Philippe Ravaud (Mar. 2018). "Misrepresentation and distortion of research in biomedical literature". In: *Proceedings of the National Academy of Sciences* 115.11, pp. 2613–2619. DOI: 10.1073/pnas.1710755115. URL: https://www.pnas.org/doi/10.1073/pnas.1710755115 (visited on 08/26/2022).

Boyd, Danah (Dec. 2007). *Why Youth (Heart) Social Network Sites: The Role of Networked Publics in Teenage Social Life*. en. SSRN Scholarly Paper. Rochester, NY. URL: https://papers.ssrn.com/abstract=1518924 (visited on 07/26/2023).

boyd, danah (Sept. 2011). "Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications". eng. In: *A Networked Self*. Rochester, NY: Routledge, pp. 47–66. ISBN: 978-0-415-80181-2. DOI: 10.4324/9780203876527-8. URL: https://papers.ssrn.com/abstract=1925128 (visited on 05/16/2023).

Brady, William J. et al. (Oct. 2019). "An ideological asymmetry in the diffusion of moralized content on social media among political leaders". eng. In: *Journal of Experimental Psychology. General* 148.10, pp. 1802–1813. ISSN: 1939-2222. DOI: 10.1037/xge0000532.

Brandes, Ulrik (1999). "Layout of graph visualizations." PhD diss. University of Konstanz. URL: http://www.ub.uni-konstanz.de/kops/volltexte/1999/255/255_1.pdf.

Brandes, Ulrik et al. (Jan. 1999). "Explorations into the Visualization of Policy Networks". en. In: *Journal of Theoretical Politics* 11.1, pp. 75–106. ISSN: 0951-6298, 1460-3667. DOI: 10.1177/

0951692899011001004. URL: http://journals.sagepub.com/doi/10.1177/0951692899011001004 (visited on 04/24/2024).

Brewer, Marilyn B. (2001). "When Does Ingroup Love Become Outgroup Hate?" In: *Social Identity, Intergroup Conflict, and Conflict Reduction*. Ed. by Richard D. Ashmore, Lee Jussim, and David Wilder. Cary, UNITED STATES: Oxford University Press, Incorporated. ISBN: 978-0-19-803143-7. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=241450 (visited on 04/16/2024).

Broniatowski, David A. et al. (Sept. 2023). "Measuring the monetization strategies of websites with application to pro- and anti-vaccine communities". en. In: *Scientific Reports* 13.1, p. 15964. ISSN: 2045-2322. DOI: 10.1038/s41598-023-43061-6. URL: https://www.nature.com/articles/s41598-023-43061-6 (visited on 05/05/2024).

Brooke, Matthew (Aug. 2022). "A Case of Twins: How Conservative Media Developed During the Golden Age of Media Objectivity". In: *American Sociological Association Annual Meeting 2022*.

Bruckman, Amy (Sept. 2002). "Studying the amateur artist: A perspective on disguising data collected in human subjects research on the Internet". en. In: *Ethics and Information Technology* 4.3, pp. 217–231. ISSN: 1572-8439. DOI: 10.1023/A:1021316409277. URL: https://doi.org/10.1023/A:1021316409277 (visited on 04/03/2020).

Bruggeman, Jeroen, V. A. Traag, and Justus Uitermark (Dec. 2012). "Detecting Communities through Network Data". en. In: *American Sociological Review* 77.6, pp. 1050–1063. ISSN: 0003-1224. DOI: 10.1177/0003122412463574. URL: https://doi.org/10.1177/0003122412463574 (visited on 09/09/2021).

Brüggemann, Michael and Sven Engesser (Jan. 2017). "Beyond false balance: How interpretive journalism shapes media coverage of climate change". en. In: *Global Environmental Change* 42, pp. 58–67. ISSN: 0959-3780. DOI: 10.1016/j.gloenvcha.2016.11.004. URL: https://www.sciencedirect.com/science/article/pii/S0959378016305209 (visited on 08/26/2022).

Bruns, Axel (2007). "Produsage". en. In: *Proceedings of the 6th ACM SIGCHI conference on Creativity & cognition - C&C '07*. Washington, DC, USA: ACM Press, p. 99. ISBN: 978-1-59593-712-4. DOI: 10.1145/1254960.1254975. URL: http://portal.acm.org/citation.cfm?doid=1254960.1254975 (visited on 06/02/2024).

– (2021). "After the 'APIcalypse': social media platforms and their fight against critical scholarly research". In: *Disinformation and Data Lockdown on Social Platforms*. Routledge. ISBN: 978-1-00-320697-2.

Bruns, Axel and Tim Highfield (May 2015). "From news blogs to news on Twitter: gatewatching and collaborative news curation". en. In: *Handbook of Digital Politics*. Edward Elgar Publishing, pp. 325–339. ISBN: 978-1-78254-876-8. URL: https://www.elgaronline.com/edcollchap/edcoll/9781782548751/9781782548751.00029.xml (visited on 04/21/2024).

Bruns, Axel and Brenda Moon (2018). "Social media in Australian Federal elections: comparing the 2013 and 2016 campaigns". In: *Journalism & Mass Communication Quarterly* 95.2, pp. 425–448.

Bruns, Axel et al. (Oct. 2017). "The Australian Twittersphere in 2016: Mapping the Follower/Followee Network". en. In: *Social Media + Society* 3.4, p. 205630511774816. ISSN: 2056-3051, 2056-3051. DOI: 10.1177/2056305117748162. URL: http://journals.sagepub.com/doi/10.1177/2056305117748162 (visited on 05/16/2024).

Burrell, Jenna and Jacob Metcalf (Apr. 2024). "Introduction for the special issue of "Ideologies of AI and the consolidation of power": Naming power". en. In: *First Monday*. ISSN: 1396-0466. DOI: 10.5210/fm.v29i4.13643. URL: https://firstmonday.org/ojs/index.php/fm/article/view/13643 (visited on 05/04/2024).

Burton, Jason W., Nicole Cruz, and Ulrike Hahn (2021). "Reconsidering evidence of moral contagion in online social networks". In: *Nature Human Behaviour* 5.12, pp. 1629–1635. URL: https://www.nature.com/articles/s41562-021-01133-5 (visited on 04/16/2024).

Butsch, Richard (2000). *The making of American audiences: from stage to television, 1750-1990*. en. Cambridge University Press. ISBN: 978-0-511-39515-4 978-0-521-66253-6 978-0-521-66483-7. URL: https://hdl.handle.net/2027/heb08300.0001.001 (visited on 05/29/2024).

– (2011). "Audiences and Publics, Media and Public Spheres". en. In: *The Handbook of Media Audiences*. John Wiley & Sons, Ltd, pp. 147–168. ISBN: 978-1-4443-4052-5. DOI: 10.1002/9781444340525.ch7. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781444340525.ch7 (visited on 05/27/2024).

Butsch, Richard and Sonia Livingstone (2014). ""Translating" Audiences, Provincializing Europe". In: *Meanings of audiences: comparative discourses*. Ed. by Richard Butsch and Sonia Livingstone, pp. 1–19. URL: https://www.academia.edu/download/35094442/Meanings_of_Audiences_-_Introduction.pdf (visited on 06/03/2024).

Cameron, Chris (Dec. 2023). "Lauren Boebert, Far-Right Firebrand, Is Switching House Districts in Colorado". en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2023/12/27/us/politics/lauren-boebert-switches-districts-colorado.html (visited on 04/08/2024).

Cameron, Michael P., Patrick Barrett, and Bob Stewardson (Oct. 2016). "Can Social Media Predict Election Results? Evidence From New Zealand". In: *Journal of Political Marketing* 15.4, pp. 416–432. ISSN: 1537-7857. DOI: 10.1080/15377857.2014.959690. URL: https://doi.org/10.1080/15377857.2014.959690 (visited on 05/08/2023).

Campan, Alina et al. (Dec. 2018). "Is Data Collection through Twitter Streaming API Useful for Academic Research?" In: *2018 IEEE International Conference on Big Data (Big Data)*, pp. 3638–3643. DOI: 10.1109/BigData.2018.8621898. URL: https://ieeexplore.ieee.org/abstract/document/8621898 (visited on 04/22/2024).

Campbell, James E. (May 1991). "The Presidential Surge and its Midterm Decline in Congressional Elections, 1868-1988". In: *The Journal of Politics* 53.2, pp. 477–487. ISSN: 0022-3816. DOI: 10.1017/S0022381600048404. URL: https://www.journals.uchicago.edu/doi/abs/10.1017/S0022381600048404 (visited on 12/03/2022).

Cann, Tristan J. B., Iain S. Weaver, and Hywel T. P. Williams (July 2020). "Is it correct to project and detect? How weighting unipartite projections influences community detection". en. In: *Network Science* 8.S1, S145–S163. ISSN: 2050-1242, 2050-1250. DOI: 10.1017/nws.2020.11. (Visited on 08/13/2022).

Carley, Kathleen M. (1999). "On the evolution of social and organizational networks". In: *Research in the Sociology of Organizations* 16.0.

Carlson, Matt (Feb. 2020). "Fake news as an informational moral panic: the symbolic deviancy of social media during the 2016 US presidential election". In: *Information, Communication & Society* 23.3, pp. 374–388. DOI: 10.1080/1369118X.2018.1505934. (Visited on 01/16/2022).

Cartwright, Dorwin and Frank Harary (1956). "Structural balance: a generalization of Heider's theory." In: *Psychological review* 63.5, p. 277.

Case, Donald O. and Georgeann M. Higgins (2000). "How can we investigate citation behavior? A study of reasons for citing literature in communication". en. In: *Journal of the American Society for Information Science* 51.7, pp. 635–645. ISSN: 1097-4571. DOI: 10.1002/(SICI)1097-4571(2000)51:7<635::AID-ASI6>3.0.CO;2-H. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291097-4571%282000%2951%3A7%3C635%3A%3AAID-ASI6%3E3.0.CO%3B2-H (visited on 10/28/2021).

Catalini, Christian, Nicola Lacetera, and Alexander Oettl (Nov. 2015). "The incidence and role of negative citations in science". In: *Proceedings of the National Academy of Sciences of the United States of America* 112.45, pp. 13823–13826. ISSN: 0027-8424. DOI: 10.1073/pnas.1502280112. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4653214/ (visited on 05/26/2022).

Center for an Informed Public et al. (2021). *The Long Fuse: Misinformation and the 2020 Election*. en. https://purl.stanford.edu/tr171zs0069. URL: https://purl.stanford.edu/tr171zs0069 (visited on 03/26/2021).

Center for Countering Digital Hate (2021). *The Disinformation Dozen*. en. https://www.counterhate.com/disinformationdozen. (Visited on 08/03/2021).

Chen, Bin, Josephine Lukito, and Gyo Hyun Koo (July 2023). "Comparing the #StopTheSteal Movement across Multiple Platforms: Differentiating Discourse on Facebook, Twitter, and Parler". en. In: *Social Media + Society* 9.3, p. 20563051231196879. ISSN: 2056-3051. DOI: 10.1177/20563051231196879. URL: https://doi.org/10.1177/20563051231196879 (visited on 05/02/2024).

Cheng, Justin et al. (2017). "Anyone can become a troll: Causes of trolling behavior in online discussions". In: *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, pp. 1217–1230.

Childs, Kiara M. (Apr. 2022). ""The Shade of It All": How Black Women Use Instagram and YouTube to Contest Colorism in the Beauty Industry". en. In: *Social Media + Society* 8.2, p. 20563051221107634. ISSN: 2056-3051. DOI: 10.1177/20563051221107634. URL: https://doi.org/10.1177/20563051221107634 (visited on 04/20/2024).

Chin, Bertha (2019). "When Hated Characters Talk Back: Twitter, Hate, and Fan/Celebrity Interactions". In: *Anti-Fandom*. Ed. by Melissa A. Click. Vol. 24. Dislike and Hate in the Digital Age. NYU Press, pp. 291–314. ISBN: 978-1-4798-0527-3. URL: https://www.jstor.org/stable/j.ctvwrm46p.17 (visited on 06/03/2024).

Chinn, Sedona, Daniel S. Lane, and Philip S. Hart (Oct. 2018). "In consensus we trust? Persuasive effects of scientific consensus communication". en. In: *Public Understanding of Science* 27.7, pp. 807–823. ISSN: 0963-6625. DOI: 10.1177/0963662518791094. URL: https://doi.org/10.1177/0963662518791094 (visited on 08/29/2022).

Choi, Sujin (2015). "The two-step flow of communication in Twitter-based public forums". In: *Social science computer review* 33.6, pp. 696–711.

Chowdhury, Farhan Asif et al. (2021). "Examining factors associated with twitter account suspension following the 2020 us presidential election". In: *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pp. 607–612.

Christin, Angèle and Rebecca Lewis (Jan. 2021). "The Drama of Metrics: Status, Spectacle, and Resistance Among YouTube Drama Creators". en. In: *Social Media + Society* 7.1, p. 205630512199966.

ISSN: 2056-3051, 2056-3051. DOI: 10.1177/2056305121999660. URL: http://journals.sagepub.com/doi/10.1177/2056305121999660 (visited on 04/15/2024).

Christin, Angèle and Yingdan Lu (Apr. 2023). "The influencer pay gap: Platform labor meets racial capitalism". en. In: *New Media & Society*, p. 14614448231164995. ISSN: 1461-4448. DOI: 10.1177/14614448231164995. URL: https://doi.org/10.1177/14614448231164995 (visited on 04/20/2024).

Chu, Derek K. et al. (June 2020). "Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: a systematic review and meta-analysis". English. In: *The Lancet* 395.10242, pp. 1973–1987. ISSN: 0140-6736, 1474-547X. DOI: 10.1016/S0140-6736(20)31142-9. URL: https://www.thelancet.com/article/S0140-6736(20)31142-9/fulltext (visited on 10/25/2022).

Chughtai, Abrar A., Holly Seale, and C. Raina Macintyre (Oct. 2020). "Effectiveness of Cloth Masks for Protection Against Severe Acute Respiratory Syndrome Coronavirus 2". eng. In: *Emerging Infectious Diseases* 26.10, e200948. ISSN: 1080-6059. DOI: 10.3201/eid2610.200948. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7510705/ (visited on 12/27/2021).

Chughtai, Abrar A. et al. (May 2020). "Policies on the use of respiratory protection for hospital health workers to protect from coronavirus disease (COVID-19)". In: *International Journal of Nursing Studies* 105, p. 103567. ISSN: 0020-7489. DOI: 10.1016/j.ijnurstu.2020.103567. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7174826/ (visited on 12/27/2021).

Cifor, Marika and Claire McDonald (Jan. 2023). ""I Hope We Leave More of a Record"Radical Queer Care within and for the AIDS INFO BBS's Caregivers Mailing List". en. In: *Feminist Media Histories* 9.1, pp. 78–97. DOI: 10.1525/fmh.2023.9.1.78. URL: https://online.ucpress.edu/fmh/article/9/1/78/195070/I-Hope-We-Leave-More-of-a-Record-Radical-Queer (visited on 11/06/2023).

Clark, Meredith D. (2019). "To tweet our own cause: A mixed-methods study of the online phenomenon "Black Twitter"". English. Ph.D. diss. United States – North Carolina: The University of North Carolina at Chapel Hill. (Visited on 10/11/2021).

– (2020). "DRAG THEM: A brief etymology of so-called "cancel culture"". In: *Communication and the Public* 5.3-4, pp. 88–92.

Click, Melissa A. (2019). "Introduction: Haters Gonna Hate". In: *Anti-Fandom*. Ed. by Melissa A. Click. Vol. 24. Dislike and Hate in the Digital Age. NYU Press, pp. 1–22. ISBN: 978-1-4798-0527-3. URL: https://www.jstor.org/stable/j.ctvwrm46p.3 (visited on 06/03/2024).

Collins, Patricia Hill (1989). "The Social Construction of Black Feminist Thought". In: *Signs* 14.4, pp. 745–773. ISSN: 0097-9740. URL: https://www.jstor.org/stable/3174683 (visited on 05/01/2024).

– (Sept. 1990). *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*. 2nd ed. New York: Routledge. ISBN: 978-0-203-90005-5. DOI: 10.4324/9780203900055.

Conforti, Costanza et al. (May 2020). *Will-They-Won't-They: A Very Large Dataset for Stance Detection on Twitter*. DOI: 10.48550/arXiv.2005.00388. URL: http://arxiv.org/abs/2005.00388 (visited on 04/03/2023).

Conrad, Ryan (2014). *Against equality: Queer revolution, not mere inclusion*. AK Press. URL: https://books.google.com/books?hl=en\&lr=\&id=2G1kCgAAQBAJ\&oi=fnd\&pg=PT3\&dq=

Against+Equality:+Queer+Revolution,+Not+Mere+Inclusion\&ots=jY8SiyabqJ\&sig=YuAF3IncooXvOaB8POMJvhRkMh0 (visited on 05/04/2024).

Conti, Mauro, Jenil Gathani, and Pier Paolo Tricomi (Aug. 2022). "Virtual Influencers in Online Social Media". In: *IEEE Communications Magazine* 60.8, pp. 86–91. ISSN: 1558-1896. DOI: 10.1109/MCOM.001.2100786. URL: https://ieeexplore.ieee.org/abstract/document/9772313 (visited on 05/07/2024).

Converse, Phillip E. (1964). "The nature of belief systems in mass publics". In: *Ideology and Discontent*.

Cook, Christine, Juliette Schaafsma, and Marjolijn Antheunis (Sept. 2018). "Under the bridge: An in-depth examination of online trolling in the gaming context". en. In: *New Media & Society* 20.9, pp. 3323–3340. ISSN: 1461-4448. DOI: 10.1177/1461444817748578. URL: https://doi.org/10.1177/1461444817748578 (visited on 11/08/2023).

Counts (Sept. 2023). "TikTok's Rules Deter Researchers From Crunching Data on Users, Misinformation". en. In: *Bloomberg*. URL: https://www.bloomberg.com/news/articles/2023-09-21/tiktok-terms-of-service-strict-for-researchers (visited on 04/16/2024).

Crampton, Jeremy W. (June 2001). "Maps as social constructions: power, communication and visualization". en. In: *Progress in Human Geography* 25.2, pp. 235–252. ISSN: 0309-1325, 1477-0288. DOI: 10.1191/030913201678580494. URL: http://journals.sagepub.com/doi/10.1191/030913201678580494 (visited on 04/29/2024).

Cross, Katherine (Nov. 2019). "Toward a formal sociology of online harassment". en. In: *Human Technology* 15.3, pp. 326–346. ISSN: 1795-6889. URL: https://ht.csr-pub.eu/index.php/ht/article/view/274 (visited on 03/24/2023).

Cushion, Stephen, Declan McDowell-Naylor, and Richard Thomas (Apr. 2021). "Why National Media Systems Matter: A Longitudinal Analysis of How UK Left-Wing and Right-Wing Alternative Media Critique Mainstream Media (2015–2018)". In: *Journalism Studies* 22.5, pp. 633–652. ISSN: 1461-670X. DOI: 10.1080/1461670X.2021.1893795. URL: https://doi.org/10.1080/1461670X.2021.1893795 (visited on 04/12/2024).

Deadspin (Jan. 2013). *realDonaldTrump Go fuck yourself.* en. Tweet. URL: https://twitter.com/Deadspin/status/291941831476920321 (visited on 05/02/2024).

Deb, Ashok et al. (May 2019). "Perils and Challenges of Social Media and Election Manipulation Analysis: The 2018 US Midterms". In: *Companion Proceedings of The 2019 World Wide Web Conference*. WWW '19. New York, NY, USA: Association for Computing Machinery, pp. 237–247. ISBN: 978-1-4503-6675-5. DOI: 10.1145/3308560.3316486. URL: https://dl.acm.org/doi/10.1145/3308560.3316486 (visited on 05/07/2023).

Derr, Tyler et al. (2019). "Balance in signed bipartite networks". In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pp. 1221–1230.

DeVito, Michael A., Darren Gergle, and Jeremy Birnholtz (May 2017). ""Algorithms ruin everything": #RIPTwitter, Folk Theories, and Resistance to Algorithmic Change in Social Media". en. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. Denver Colorado USA: ACM, pp. 3163–3174. ISBN: 978-1-4503-4655-9. DOI: 10.1145/3025453.3025659. URL: https://dl.acm.org/doi/10.1145/3025453.3025659 (visited on 06/04/2024).

Devlin, Jacob et al. (May 2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. URL: http://arxiv.org/abs/1810.04805 (visited on 03/03/2024).

Dianati, Navid (Jan. 2016). "Unwinding the hairball graph: Pruning algorithms for weighted complex networks". In: *Physical Review E* 93.1, p. 012304. DOI: 10.1103/PhysRevE.93.012304. (Visited on 08/30/2021).

Disease Control & Prevention, Centers for (2022). "Masks". In: URL: https://www.cdc.gov/coronavirus/2019-ncov/prevent-getting-sick/masks.html.

Donovan, Joan and danah boyd (Feb. 2021). "Stop the Presses? Moving From Strategic Silence to Strategic Amplification in a Networked Media Ecosystem". en. In: *American Behavioral Scientist* 65.2, pp. 333–350. ISSN: 0002-7642. DOI: 10.1177/0002764219878229. URL: https://doi.org/10.1177/0002764219878229 (visited on 03/24/2023).

Donovan, Joan, Kaylee Fagan, and Frances Lee (July 2022). *"President Trump is Calling Us to Fight:" What the Court Documents Reveal About the Motivations Behind January 6 and Networked Incitement*. en. URL: https://mediamanipulation.org/research/president-trump-calling-us-fight-what-court-documents-reveal-about-motivations-behind (visited on 01/05/2024).

Doreian, Patrick (1988). "Using multiple network analytic tools for a single social network". en. In: *Social Networks* 10.4, pp. 287–312. DOI: 10.1016/0378-8733(88)90001-9. (Visited on 07/08/2021).

Downey, John and Natalie Fenton (2003). "New media, counter publicity and the public sphere". In: *New media & society* 5.2, pp. 185–202.

Driel, Loes van and Delia Dumitrica (Feb. 2021). "Selling brands while staying "Authentic": The professionalization of Instagram influencers". en. In: *Convergence* 27.1, pp. 66–84. ISSN: 1354-8565. DOI: 10.1177/1354856520902136. URL: https://doi.org/10.1177/1354856520902136 (visited on 05/07/2024).

Duffett, Mark (2013). *Understanding Fandom: An Introduction to the Study of Media Fan Culture*. New York, UNITED STATES: Bloomsbury Publishing. ISBN: 978-1-62356-585-5. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=1363769 (visited on 06/03/2024).

Duffy, Brooke Erin (2017). *(Not) getting paid to do what you love: Gender, social media, and aspirational work*. Yale University Press. URL: https://books.google.com/books?hl=en\&lr=\&id=s0wmDwAAQBAJ\&oi=fnd\&pg=PP1\&dq=info:PUnru9-Xyd0J:scholar.google.com\&ots=bXHgICYWac\&sig=1Z20YlrIciLYUWc_ceQ2cjXogMQ (visited on 04/15/2024).

Duffy, Brooke Erin and Colten Meisner (Mar. 2023). "Platform governance at the margins: Social media creators' experiences with algorithmic (in)visibility". en. In: *Media, Culture & Society* 45.2, pp. 285–304. ISSN: 0163-4437. DOI: 10.1177/01634437221111923. URL: https://doi.org/10.1177/01634437221111923 (visited on 04/21/2024).

Duffy, Brooke Erin, Kate M. Miltner, and Amanda Wahlstedt (July 2022). "Policing "Fake" Femininity: Authenticity, Accountability, and Influencer Antifandom". en. In: *New Media & Society* 24.7, pp. 1657–1676. ISSN: 1461-4448. DOI: 10.1177/14614448221099234. URL: https://doi.org/10.1177/14614448221099234 (visited on 04/20/2024).

Duffy, Brooke Erin and Elizabeth Wissinger (Oct. 2017). "Mythologies of Creative Work in the Social Media Age: Fun, Free, and "Just Being Me"". en. In: *International Journal of Communication* 11.0, p. 20. ISSN: 1932-8036. URL: https://ijoc.org/index.php/ijoc/article/view/7322 (visited on 05/07/2024).

Duffy, Brooke Erin et al. (Apr. 2021). "The Nested Precarities of Creative Labor on Social Media". en. In: *Social Media + Society* 7.2, p. 20563051211021368. ISSN: 2056-3051. DOI: 10.1177/20563051211021368. URL: https://doi.org/10.1177/20563051211021368 (visited on 04/20/2024).

Duffy, Clare (Nov. 2022). "Elon Musk unbans several controversial Twitter accounts, but says not yet on Trump | CNN Business". en. In: *CNN*. URL: https://www.cnn.com/2022/11/18/tech/elon-musk-twitter-content-policy/index.html (visited on 11/03/2023).

Duggan, Maeve (May 2017). *How Platforms Are Poisoning Conversations*. en. URL: https://www.theatlantic.com/technology/archive/2017/05/how-platforms-are-poisoning-conversations/524031/ (visited on 05/01/2024).

Dusenbery, Maya (Mar. 2018). *Doing Harm: The Truth About How Bad Medicine and Lazy Science Leave Women Dismissed, Misdiagnosed, and Sick*. en. HarperCollins. ISBN: 978-0-06-247081-2.

Dym, Brianna and Casey Fiesler (June 2020). "Ethical and privacy considerations for research using online fandom data". en. In: *Transformative Works and Cultures* 33. ISSN: 1941-2258. DOI: 10.3983/twc.2020.1733. URL: https://journal.transformativeworks.org/index.php/twc/article/view/1733 (visited on 04/21/2024).

Ecker, Ullrich K. H., Stephan Lewandowsky, and Matthew Chadwick (Aug. 2020). "Can corrections spread misinformation to new audiences? Testing for the elusive familiarity backfire effect". In: *Cognitive Research: Principles and Implications* 5.1, p. 41. ISSN: 2365-7464. DOI: 10.1186/s41235-020-00241-6. URL: https://doi.org/10.1186/s41235-020-00241-6 (visited on 05/24/2023).

Eckert, Penelope and Sally McConnell-Ginet (1992). "Think practically and look locally: Language and gender as community-based practice". In: *Annual Review of Anthropology* 21, pp. 461–490. ISSN: 1545-4290. DOI: 10.1146/annurev.an.21.100192.002333.

Edge, David (1979). "Quantitative measures of communication in science: A critical review". In: *History of science* 17.2, pp. 102–134.

Edsall, Thomas B. (Feb. 2021). "Opinion | Democracy Is Weakening Right in Front of Us". en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2021/02/17/opinion/digital-revolution-democracy-fake-news.html (visited on 05/01/2024).

– (June 2022). "Opinion | We're Staring at Our Phones, Full of Rage for 'the Other Side'". en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2022/06/15/opinion/social-media-polarization-democracy.html (visited on 02/22/2024).

Edwards, Paul N et al. (2013). "Knowledge infrastructures: Intellectual frameworks and research challenges". In.

Fang, Jenn (May 2021). *Social media sites popular with Asian Americans have a big misinformation problem*. en-US. http://prismreports.org/2021/05/26/social-media-sites-often-used-by-asian-americans-have-a-big-problem-with-right-wing-misinformation/. (Visited on 09/27/2021).

Farrell, Henry and Daniel W. Drezner (Jan. 2008). "The power and politics of blogs". en. In: *Public Choice* 134.1, pp. 15–30. ISSN: 1573-7101. DOI: 10.1007/s11127-007-9198-1. URL: https://doi.org/10.1007/s11127-007-9198-1 (visited on 04/21/2024).

Feagin, Joe R., Anthony M. Orum, and Gideon Sjoberg (1991). *A Case for the Case Study*. UNC Press Books. URL: https://books.google.com/books?hl=en\&lr=\&id=7A39B6ZLyJQC\&oi=fnd\&pg=PA1\&ots=H15y_4hjc1\&sig=cn1L-Cu7PiS-Yl99IIHlbe48ZPw (visited on 04/29/2024).

Feder, Shira (Sept. 2021). "What to do when someone you care about keeps getting into fights online". en-US. In: *Washington Post*. ISSN: 0190-8286. URL: https://www.washingtonpost.com/lifestyle/wellness/online-argument-mental-health-help/2021/09/24/0e60d8fc-1c8e-11ec-bcb8-0cb135811007_story.html (visited on 05/01/2024).

Feezell, Jessica T. (June 2018). "Agenda Setting through Social Media: The Importance of Incidental News Exposure and Social Filtering in the Digital Era". en. In: *Political Research Quarterly* 71.2, pp. 482–494. ISSN: 1065-9129. DOI: 10.1177/1065912917744895. URL: https://doi.org/10.1177/1065912917744895 (visited on 10/20/2023).

Fehr, Ernst and Simon Gächter (Jan. 2002). "Altruistic punishment in humans". en. In: *Nature* 415.6868, pp. 137–140. ISSN: 1476-4687. DOI: 10.1038/415137a. URL: https://www.nature.com/articles/415137a (visited on 04/16/2024).

Fiesler, Casey, Cliff Lampe, and Amy S. Bruckman (Feb. 2016). "Reality and Perception of Copyright Terms of Service for Online Content Creation". In: *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. CSCW '16. San Francisco, California, USA: Association for Computing Machinery, pp. 1450–1461. ISBN: 978-1-4503-3592-8. DOI: 10.1145/2818048.2819931. URL: https://doi.org/10.1145/2818048.2819931 (visited on 04/23/2020).

Fiesler, Casey and Nicholas Proferes (Jan. 2018). ""Participant" Perceptions of Twitter Research Ethics". en. In: *Social Media + Society* 4.1, p. 2056305118763366. ISSN: 2056-3051. DOI: 10.1177/2056305118763366. URL: https://doi.org/10.1177/2056305118763366 (visited on 04/03/2020).

Fischhoff, Baruch and Dietram A. Scheufele (Aug. 2013). "The science of science communication". In: *Proceedings of the National Academy of Sciences* 110.supplement_3, pp. 14031–14032. DOI: 10.1073/pnas.1312080110. URL: https://www.pnas.org/doi/full/10.1073/pnas.1312080110 (visited on 12/18/2022).

Fiske, Susan T. and Cydney Dupree (Sept. 2014). "Gaining trust as well as respect in communicating to motivated audiences about science topics". In: *Proceedings of the National Academy of Sciences* 111.supplement_4, pp. 13593–13597. DOI: 10.1073/pnas.1317505111. URL: https://www.pnas.org/doi/full/10.1073/pnas.1317505111 (visited on 12/18/2022).

Flamino, James et al. (Mar. 2023). "Political polarization of news media and influencers on Twitter in the 2016 and 2020 US presidential elections". en. In: *Nature Human Behaviour*, pp. 1–13. ISSN: 2397-3374. DOI: 10.1038/s41562-023-01550-8. URL: https://www.nature.com/articles/s41562-023-01550-8 (visited on 05/23/2023).

Fortuna, Paula, Juan Soler, and Leo Wanner (2020). "Toxic, hateful, offensive or abusive? what are we really classifying? an empirical analysis of hate speech datasets". In: *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pp. 6786–6794. URL: https://aclanthology.org/2020.lrec-1.838/ (visited on 04/16/2024).

Fortunato, Santo (2010). "Community detection in graphs". In: *Physics reports* 486.3-5, pp. 75–174. URL: https://www.sciencedirect.com/science/article/pii/S0370157309002841 (visited on 04/24/2024).

Fortunato, Santo et al. (2018). "Science of science". In: *Science* 359.6379, eaao0185.

Foucault Welles, Brooke and Isabel Meirelles (Feb. 2015). "Visualizing Computational Social Science: The Multiple Lives of a Complex Image". en. In: *Science Communication* 37.1, pp. 34–58. DOI: 10.1177/1075547014556540. (Visited on 01/16/2022).

franzke, aline shakti et al. (2020). "Internet research: Ethical guidelines 3.0". In: URL: https://pure.au.dk/portal/en/publications/internet-research-ethical-guidelines-30 (visited on 05/01/2024).

Fraser, Nancy (1990). "Rethinking the Public Sphere: A Contribution to the Critique of Actually Existing Democracy". In: *Social Text* 25/26, pp. 56–80. ISSN: 0164-2472. DOI: 10.2307/466240. URL: https://www.jstor.org/stable/466240 (visited on 12/14/2022).

Freelon, Deen, Charlton McIlwain, and Meredith Clark (2018). "Quantifying the power and consequences of social media protest". In: *new media & society* 20.3, pp. 990–1011.

Freelon, Deen, Charlton D. McIlwain, and Meredith Clark (Feb. 2016). *Beyond the Hashtags: #Ferguson, #Blacklivesmatter, and the Online Struggle for Offline Justice*. en. SSRN Scholarly Paper ID 2747066. Rochester, NY: Social Science Research Network. DOI: 10.2139/ssrn.2747066. URL: https://papers.ssrn.com/abstract=2747066 (visited on 08/30/2021).

Gallagher, Erin (Mar. 2020). *Trump Trains*. en. https://erin-gallagher.medium.com/trump-trains-84bea1c3170d. URL: https://erin-gallagher.medium.com/trump-trains-84bea1c3170d (visited on 08/24/2021).

Gallagher, Ryan J. et al. (Apr. 2021). "Sustained Online Amplification of COVID-19 Elites in the United States". en. In: *Social Media + Society* 7.2, p. 20563051211024957. ISSN: 2056-3051. DOI: 10.1177/20563051211024957. URL: https://doi.org/10.1177/20563051211024957 (visited on 01/16/2022).

Gannon, Valerie and Andrea Prothero (2018). "Beauty bloggers and YouTubers as a community of practice". In: *Journal of Marketing Management* 34.7-8, pp. 592–619. URL: https://www.tandfonline.com/doi/abs/10.1080/0267257X.2018.1482941 (visited on 11/06/2023).

Garcia, David, Pavlin Mavrodiev, and Frank Schweitzer (Oct. 2013). "Social resilience in online communities: the autopsy of friendster". In: *Proceedings of the first ACM conference on Online social networks*. COSN '13. New York, NY, USA: Association for Computing Machinery, pp. 39–50. ISBN: 978-1-4503-2084-9. DOI: 10.1145/2512938.2512946. URL: https://dl.acm.org/doi/10.1145/2512938.2512946 (visited on 05/08/2023).

Garimella, Kiran, Ingmar Weber, and Munmun De Choudhury (May 2016). "Quote RTs on Twitter: usage of the new feature for political discourse". In: *Proceedings of the 8th ACM Conference on Web Science*. WebSci '16. New York, NY, USA: Association for Computing Machinery, pp. 200–204. ISBN: 978-1-4503-4208-7. DOI: 10.1145/2908131.2908170. URL: https://doi.org/10.1145/2908131.2908170 (visited on 04/12/2024).

Gayo-Avello, Daniel (Dec. 2013). "A Meta-Analysis of State-of-the-Art Electoral Prediction From Twitter Data". en. In: *Social Science Computer Review* 31.6, pp. 649–679. ISSN: 0894-4393. DOI: 10.1177/0894439313493979. URL: https://doi.org/10.1177/0894439313493979 (visited on 05/08/2023).

Geiger, R. Stuart and David Ribes (2011). "Trace ethnography: Following coordination through documentary practices". In: *2011 44th Hawaii international conference on system sciences*. IEEE, pp. 1–10. URL: https://ieeexplore.ieee.org/abstract/document/5718606 (visited on 04/29/2024).

Gerring, John (2012). "Mere description". In: *British Journal of Political Science* 42.4, pp. 721–746. URL: https://www.cambridge.org/core/journals/british-journal-of-political-science/article/mere-description/833643C6242D3A45D48BAAC3EF0C33D0 (visited on 04/30/2024).

Ghoshal, Gourab, Liping Chi, and Albert-László Barabási (2013). "Uncovering the role of elementary processes in network evolution". In: *Scientific reports* 3.1, pp. 1–8.

Giddens, Anthony (1984). *The constitution of society: Outline of the theory of structuration*. Univ of California Press. URL: https://books.google.com/books?hl=en\&lr=\&id=x2bf4g9Z6ZwC\&oi=fnd\&pg=PR9\&dq=The+Constitution+of+Society+giddens\&ots=jPUP3rzv6y\&sig=Q_-xH3RQEJfmNmZ4uMzppC2ZGDg (visited on 06/03/2024).

Gieryn, Thomas F. (1983). "Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists". In: *American Sociological Review* 48.6, pp. 781–795. ISSN: 0003-1224. DOI: 10.2307/2095325. URL: https://www.jstor.org/stable/2095325 (visited on 12/11/2020).

Gillespie, Tarleton (June 2023). "The Fact of Content Moderation; Or, Let's Not Solve the Platforms' Problems for Them". en. In: *Media and Communication* 11.2, pp. 406–409. ISSN: 2183-2439. URL: https://www.cogitatiopress.com/mediaandcommunication/article/view/6610 (visited on 05/04/2024).

Gitlin, Todd (1978). "Media Sociology: The Dominant Paradigm". In: *Theory and Society* 6.2, pp. 205–253. ISSN: 0304-2421. URL: https://www.jstor.org/stable/657009 (visited on 12/12/2022).

Godoy, Maria (June 2020). *Yes, Wearing Masks Helps. Here's Why*. en. URL: https://www.npr.org/sections/health-shots/2020/06/21/880832213/yes-wearing-masks-helps-heres-why (visited on 11/02/2022).

Goffman, Erving (1959). *The Presentation of Self in Everyday Life*. 3rd ed. Doubleday. ISBN: 978-0-14-013571-8.

Goggin, Ben (Mar. 2024). "Big Tech companies reveal trust and safety cuts in disclosures to Senate Judiciary Committee". en. In: *NBC News*. URL: https://www.nbcnews.com/tech/tech-news/big-tech-companies-reveal-trust-safety-cuts-disclosures-senate-judicia-rcna145435 (visited on 04/16/2024).

González-Bailón, Sandra and Yphtach Lelkes (2023). "Do social media undermine social cohesion? A critical review". en. In: *Social Issues and Policy Review* 17.1, pp. 155–180. ISSN: 1751-2409. DOI: 10.1111/sipr.12091. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/sipr.12091 (visited on 04/16/2024).

Goodman, Amy (Apr. 2016). *Democracy Now!: Twenty Years Covering the Movements Changing America*. en. Simon and Schuster. ISBN: 978-1-5011-2360-3.

Gottfried, Jeffrey (Jan. 2024). *Americans' Social Media Use*. en-US. URL: https://www.pewresearch.org/internet/2024/01/31/americans-social-media-use/ (visited on 04/16/2024).

Graham, Sage L. and Claire Hardaker (2017). "(Im)politeness in Digital Communication". en. In: *The Palgrave Handbook of Linguistic (Im)politeness*. Ed. by Jonathan Culpeper, Michael Haugh, and Dániel Z. Kádár. London: Palgrave Macmillan UK, pp. 785–814. ISBN: 978-1-137-37507-0 978-1-137-37508-7. DOI: 10.1057/978-1-137-37508-7_30. URL: http://link.springer.com/10.1057/978-1-137-37508-7_30 (visited on 04/16/2024).

Gray, Jonathan (June 2021). *Dislike-Minded: Media, Audiences, and the Dynamics of Taste*. New York University Press. ISBN: 978-1-4798-0999-8. DOI: 10.18574/nyu/9781479809998.001.0001. URL: https://www.degruyter.com/document/doi/10.18574/nyu/9781479809998.001.0001/html (visited on 04/12/2024).

Gray, Kishonna L., Bertan Buyukozturk, and Zachary G. Hill (Mar. 2017). "Blurring the boundaries: Using Gamergate to examine "real" and symbolic violence against women in contemporary gaming culture". en. In: *Sociology Compass* 11.3, e12458. ISSN: 1751-9020, 1751-9020. DOI: 10.1111/soc4.12458. URL: https://compass.onlinelibrary.wiley.com/doi/10.1111/soc4.12458 (visited on 04/12/2024).

Gray, Kishonna L. and Krysten Stein (Aug. 2021). ""We 'said her name' and got zucked": Black Women Calling-out the Carceral Logics of Digital Platforms". en. In: *Gender & Society* 35.4, pp. 538–545. ISSN: 0891-2432. DOI: 10.1177/08912432211029393. URL: https://doi.org/10.1177/08912432211029393 (visited on 04/21/2024).

Gray, Rosie (Feb. 2024). *Donald Trump Jr. Ventures Into the Lifestyle Space – and Brings the Culture War With Him*. en. URL: https://www.politico.com/news/magazine/2024/02/16/donald-trump-jr-lifestyle-influencer-magazine-00131445 (visited on 04/08/2024).

Green, Joshua and Henry Jenkins (Apr. 2011). "Spreadable Media: How Audiences Create Value and Meaning in a Networked Economy". en. In: pp. 109–127. DOI: 10.1002/9781444340525.ch5. URL: https://onlinelibrary.wiley.com/doi/10.1002/9781444340525.ch5 (visited on 05/28/2024).

Gregorio, Giovanni De and Catalina Goanta (Mar. 2022). "The Influencer Republic: Monetizing Political Speech on Social Media". en. In: *German Law Journal* 23.2, pp. 204–225. ISSN: 2071-8322. DOI: 10.1017/glj.2022.15. URL: https://www.cambridge.org/core/journals/german-law-journal/article/influencer-republic-monetizing-political-speech-on-social-media/C96153D8918E8F5B2C0E24EF233678D5 (visited on 05/05/2024).

Grootendorst, Maarten (Mar. 2022). *BERTopic: Neural topic modeling with a class-based TF-IDF procedure*. URL: http://arxiv.org/abs/2203.05794 (visited on 03/19/2024).

Guan, rank (Sept. 2019). "The Numbered World". en-US. In: *The Baffler*. URL: https://thebaffler.com/outbursts/the-numbered-world-guan (visited on 07/30/2023).

Gummesson, Evert (2007). "Case study research and network theory: birds of a feather". In: *Qualitative Research in Organizations and Management: An International Journal* 2.3, pp. 226–248. URL: https://www.emerald.com/insight/content/doi/10.1108/17465640710835373/full/html (visited on 04/30/2024).

Gutiérrez, Arcelia (Jan. 2022). "Situating Representation As a Form of Erasure: #OscarsSoWhite, Black Twitter, and Latinx Twitter". en. In: *Television & New Media* 23.1, pp. 100–118. ISSN: 1527-4764, 1552-8316. DOI: 10.1177/1527476420961247. URL: http://journals.sagepub.com/doi/10.1177/1527476420961247 (visited on 04/17/2024).

Habermas, Jurgen (1991). *The structural transformation of the public sphere: An inquiry into a category of bourgeois society*. MIT press.

Habermas, Jürgen (2006). "Political Communication in Media Society: Does Democracy Still Enjoy an Epistemic Dimension? The Impact of Normative Theory on Empirical Research1". en. In: *Communication Theory* 16.4, pp. 411–426. ISSN: 1468-2885. DOI: 10.1111/j.1468-2885.2006.00280.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-2885.2006.00280.x (visited on 05/18/2023).

Haimson, Oliver L. et al. (Oct. 2021). "Disproportionate Removals and Differing Content Moderation Experiences for Conservative, Transgender, and Black Social Media Users: Marginalization and Moderation Gray Areas". In: *Proceedings of the ACM on Human-Computer Interaction*

5.CSCW2, 466:1–466:35. DOI: 10.1145/3479610. URL: https://dl.acm.org/doi/10.1145/3479610 (visited on 05/16/2024).

Hall, Stuart (1980). "Encoding/Decoding". In: *Media Studies*. Ed. by Sue Thornham, Caroline Bassett, and Paul Marris. A Reader. Edinburgh University Press, pp. 28–38. ISBN: 978-0-7486-3783-6. URL: https://www.jstor.org/stable/10.3366/j.ctvxcrv1h.7 (visited on 06/02/2024).

– (1982). "The rediscovery of 'ideology': Return of the repressed in media studies". In: *Culture, society and the media*. Routledge, pp. 61–95.

Ham, F. Van and M. Wattenberg (2008). "Centrality Based Visualization of Small World Graphs". en. In: *Computer Graphics Forum* 27.3, pp. 975–982. DOI: 10.1111/j.1467-8659.2008.01232.x. (Visited on 08/30/2021).

Hamilton, Amber M. (July 2020). "A Genealogy of Critical Race and Digital Studies: Past, Present, and Future". en. In: *Sociology of Race and Ethnicity* 6.3, pp. 292–301. ISSN: 2332-6492. DOI: 10.1177/2332649220922577. URL: https://doi.org/10.1177/2332649220922577 (visited on 05/04/2024).

Hampton, Cynthia, David Reeping, and Desen Sevi Ozkan (2021). "Positionality statements in engineering education research: A look at the hand that guides the methodological tools". In: *Studies in Engineering Education* 1.2, pp. 126–141. URL: https://pdfs.semanticscholar.org/c0e4/30440c2532c237f50c143f55e08442091a97.pdf (visited on 04/12/2024).

Han, Catherine et al. (Apr. 2023). "Hate Raids on Twitch: Echoes of the Past, New Modalities, and Implications for Platform Governance". In: *Proceedings of the ACM on Human-Computer Interaction* 7.CSCW1, 133:1–133:28. DOI: 10.1145/3579609. URL: https://dl.acm.org/doi/10.1145/3579609 (visited on 04/21/2024).

Hanson, Joe (July 2020). *How Well Do Masks Work? (Schlieren Imaging In Slow Motion!)* URL: https://www.youtube.com/watch?v=0Tp0zB904Mc (visited on 11/17/2022).

Haraway, Donna (1988). "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective". In: *Feminist Studies* 14.3, pp. 575–599. ISSN: 0046-3663. DOI: 10.2307/3178066. URL: http://www.jstor.org/stable/3178066 (visited on 12/15/2021).

Harding, Sandra (1991). *Whose science? Whose knowledge?: Thinking from women's lives*. Cornell University Press. URL: https://books.google.com/books?hl=en\&lr=\&id=eSmPEH7-u2oC\&oi=fnd\&pg=PR7\&dq=sandra+harding\&ots=Yfx_DMZkF_\&sig=-Tlf1kliI1brkNdJj3bkLIQpgts (visited on 05/01/2024).

Hargittai, Eszter (Feb. 2020). "Potential Biases in Big Data: Omitted Voices on Social Media". en. In: *Social Science Computer Review* 38.1, pp. 10–24. ISSN: 0894-4393. DOI: 10.1177/0894439318788322. URL: https://doi.org/10.1177/0894439318788322 (visited on 05/04/2024).

Harris, Brandon C., Maxwell Foxman, and William C. Partin (Apr. 2023). ""Don't Make Me Ratio You Again": How Political Influencers Encourage Platformed Political Participation". en. In: *Social Media + Society* 9.2, p. 20563051231177944. ISSN: 2056-3051. DOI: 10.1177/20563051231177944. URL: https://doi.org/10.1177/20563051231177944 (visited on 11/06/2023).

Harry, Sydette (Feb. 2021). *Social media was supposed to give everyone a voice. It didn't.* en-US. URL: https://www.marketplace.org/shows/marketplace-tech/black-women-internet-harassment-abuse-social-media-was-supposed-to-give-everyone-a-voice-it-didnt/ (visited on 03/15/2023).

Hayes, Andrew F. and Klaus Krippendorff (Apr. 2007). "Answering the Call for a Standard Reliability Measure for Coding Data". en. In: *Communication Methods and Measures* 1.1, pp. 77–89. ISSN: 1931-2458, 1931-2466. DOI: 10.1080/19312450709336664. URL: http://www.tandfonline.com/doi/abs/10.1080/19312450709336664 (visited on 03/19/2024).

Hearn, Alison and Sarah Banet-Weiser (Jan. 2020). "The Beguiling: Glamour in/as Platformed Cultural Production". en. In: *Social Media + Society* 6.1, p. 2056305119898779. ISSN: 2056-3051. DOI: 10.1177/2056305119898779. URL: https://doi.org/10.1177/2056305119898779 (visited on 05/05/2024).

Hemmer, Nicole (Aug. 2022). *Partisans: The Conservative Revolutionaries Who Remade American Politics in the 1990s*. en. Basic Books. ISBN: 978-1-5416-4687-2.

Herasimenka, Aliaksandr et al. (Dec. 2022). "The political economy of digital profiteering: communication resource mobilization by anti-vaccination actors". In: *The Journal of Communication* 73.2, pp. 126–137. ISSN: 0021-9916. DOI: 10.1093/joc/jqac043. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10066223/ (visited on 05/05/2024).

Herman, Edward S. and Noam Chomsky (1988). *Manufacturing Consent: The Political Economy of the Mass Media*. en. Pantheon Books. ISBN: 978-0-679-72034-8.

Herring, S.C. et al. (Jan. 2005). "Conversations in the Blogosphere: An Analysis "From the Bottom Up"". In: *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*, 107b–107b. DOI: 10.1109/HICSS.2005.167. URL: https://ieeexplore.ieee.org/abstract/document/1385453 (visited on 04/21/2024).

Hilbert, Martin et al. (2017). "One Step, Two Step, Network Step? Complementary Perspectives on Communication Flows in Twittered Citizen Protests". en. In: *SOCIAL SCIENCE COMPUTER REVIEW* 35.4, pp. 444–461. ISSN: 0894-4393. DOI: 10.1177/0894439316639561. URL: https://escholarship.org/uc/item/0nn4p7mv (visited on 12/12/2022).

Hitlin, Paul (Dec. 2016). "Health issues topped the list of scientific studies reaching wide audiences in 2016". en-US. In: *Pew Research Center*. URL: https://www.pewresearch.org/short-reads/2016/12/28/health-issues-topped-the-list-of-scientific-studies-reaching-wide-audiences-in-2016/ (visited on 05/25/2023).

Hogarth, Rana A. (2017). *Medicalizing Blackness: making racial difference in the Atlantic world, 1780-1840*. UNC Press Books.

Holt, Kristoffer, Tine Ustad Figenschou, and Lena Frischlich (Aug. 2019). "Key Dimensions of Alternative News Media". In: *Digital Journalism* 7.7, pp. 860–869. ISSN: 2167-0811. DOI: 10.1080/21670811.2019.1625715. URL: https://doi.org/10.1080/21670811.2019.1625715 (visited on 07/29/2022).

Holton, Avery E and Logan Molyneux (Feb. 2017). "Identity lost? The personal impact of brand journalism". en. In: *Journalism* 18.2, pp. 195–210. ISSN: 1464-8849. DOI: 10.1177/1464884915608816. URL: https://doi.org/10.1177/1464884915608816 (visited on 05/02/2024).

hooks, bell (1989). "Choosing the Margin as a Space of Radical Openness". In: *Framework: The Journal of Cinema and Media* 36, pp. 15–23. URL: https://www.ias.edu/sites/default/files/sss/pdfs/Crisis-and-Critique-2018-19/hooks_choosing_the_margins.pdf (visited on 05/01/2024).

hooks, bell (1992). "The oppositional gaze: Black female spectators". In: *Black Looks*. South End Press. URL: https://api.taylorfrancis.com/content/chapters/edit/download?identifierName=doi\&identifierValue=10.4324/9780203873304-22\&type=chapterpdf (visited on 05/31/2024).

Howard, Jeremy et al. (Apr. 2020). "Face masks against COVID-19: an evidence review". In: *Preprints.org*. URL: https://files.fast.ai/papers/masks_lit_review.pdf.

Huber, Colleen (July 2020). *Masks are neither effective nor safe*. en. URL: https://www.primarydoctor.org/masks-not-effect (visited on 12/01/2022).

Hudson, James M. and Amy Bruckman (Apr. 2004). ""Go Away": Participant Objections to Being Studied and the Ethics of Chatroom Research". In: *The Information Society* 20.2, pp. 127–139. ISSN: 0197-2243. DOI: 10.1080/01972240490423030. URL: https://doi.org/10.1080/01972240490423030 (visited on 04/10/2020).

Hund, Emily (2023). *The Influencer Industry: The Quest for Authenticity on Social Media*. Princeton University Press. URL: https://books.google.com/books?hl=en\&lr=\&id=HOOCEAAAQBAJ\&oi=fnd\&pg=PR9\&dq=influencer+industry\&ots=GEVnvlqTzc\&sig=zeWYlFJ2Fu2UQ41k_ULO6VXZkV4 (visited on 11/07/2023).

Iriberri, Alicia and Gondy Leroy (Feb. 2009). "A life-cycle perspective on online community success". In: *ACM Computing Surveys* 41.2, 11:1–11:29. ISSN: 0360-0300. DOI: 10.1145/1459352.1459356. URL: https://dl.acm.org/doi/10.1145/1459352.1459356 (visited on 05/09/2023).

Ivanova, Irina (July 2023). "Twitter is now X. Here's what that means." en-US. In: *CBS News*. URL: https://www.cbsnews.com/news/twitter-rebrand-x-name-change-elon-musk-what-it-means/ (visited on 05/01/2024).

Jack, Caroline (2017). "Lexicon of lies: Terms for problematic information". In: *Data & Society* 3.22, pp. 1094–1096.

Jack Posobiec (2022). en. Tech. rep. Southern Poverty Law Center. URL: https://www.splcenter.org/fighting-hate/extremist-files/individual/jack-posobiec (visited on 04/08/2024).

Jackson, Sarah J., Moya Bailey, and Brooke Foucault Welles (2020). # *HashtagActivism: Networks of race and gender justice*. Mit Press.

Jackson, Sarah J. and Brooke Foucault Welles (Dec. 2015). "Hijacking #MYNYPD: Social Media Dissent and Networked Counterpublics". In: *Journal of Communication* 65.6, pp. 932–952. ISSN: 0021-9916. DOI: 10.1111/jcom.12185. URL: https://doi.org/10.1111/jcom.12185 (visited on 08/30/2021).

– (2016). "# Ferguson is everywhere: Initiators in emerging counterpublic networks". In: *Information, Communication & Society* 19.3, pp. 397–418.

Jackson, Sarah J and Daniel Kreiss (Aug. 2023). "Recentering power: conceptualizing counterpublics and defensive publics". In: *Communication Theory* 33.2-3, pp. 102–111. ISSN: 1468-2885. DOI: 10.1093/ct/qtad004. URL: https://doi.org/10.1093/ct/qtad004 (visited on 05/02/2024).

Jacobin – Bias and Credibility (Apr. 2024). en-US. URL: https://mediabiasfactcheck.com/jacobin/ (visited on 04/09/2024).

Jacobs, Joshua L. et al. (June 2009). "Use of surgical face masks to reduce the incidence of the common cold among health care workers in Japan: a randomized controlled trial". eng. In: *American Journal of Infection Control* 37.5, pp. 417–419. ISSN: 1527-3296. DOI: 10.1016/j.ajic.2008.11.002.

Jacomy, Mathieu (2021). "Situating Visual Network Analysis". en-US. Billet. Aalborg University. URL: https://reticular.hypotheses.org/1879 (visited on 10/04/2022).

Jacomy, Mathieu et al. (2014). "ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software". eng. In: *PloS one* 9.6, e98679. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0098679.

Jaidka, Kokil, Subhayan Mukerjee, and Yphtach Lelkes (Apr. 2023). "Silenced on social media: the gatekeeping functions of shadowbans in the American Twitterverse". In: *Journal of Communication* 73.2, pp. 163–178. ISSN: 0021-9916. DOI: 10.1093/joc/jqac050. URL: https://doi.org/10.1093/joc/jqac050 (visited on 06/03/2024).

Jamieson, Kathleen Hall et al. (Aug. 2017). "The Political Uses and Abuses of Civility and Incivility". In: *The Oxford Handbook of Political Communication*. Ed. by Kate Kenski and Kathleen Hall Jamieson. Oxford University Press, p. 0. ISBN: 978-0-19-979347-1. DOI: 10.1093/oxfordhb/9780199793471.013.79_update_001. URL: https://doi.org/10.1093/oxfordhb/9780199793471.013.79_update_001 (visited on 04/16/2024).

Jaramillo-Dent, Daniela, Paloma Contreras-Pulido, and Amor Pérez-Rodríguez (Feb. 2022). "Immigrant Influencers on TikTok: Diverse Microcelebrity Profiles and Algorithmic (In)Visibility". en. In: *Media and Communication* 10.1, pp. 208–221. ISSN: 2183-2439. URL: https://www.cogitatiopress.com/mediaandcommunication/article/view/4743 (visited on 04/20/2024).

Javed, Muhammad Aqib et al. (Apr. 2018). "Community detection in networks: A multidisciplinary review". In: *Journal of Network and Computer Applications* 108, pp. 87–111. ISSN: 1084-8045. DOI: 10.1016/j.jnca.2018.02.011. URL: https://www.sciencedirect.com/science/article/pii/S1084804518300560 (visited on 04/24/2024).

Javid, Babak, Michael P. Weekes, and Nicholas J. Matheson (Apr. 2020). "Covid-19: should the public wear face masks?" en. In: *BMJ* 369, p. m1442. ISSN: 1756-1833. DOI: 10.1136/bmj.m1442. URL: https://www.bmj.com/content/369/bmj.m1442 (visited on 12/02/2021).

jcho710 (June 2020). *"An N-95 mask filters..."* en. Tweet. URL: https://twitter.com/jcho710/status/1273924949162340353 (visited on 06/10/2023).

Jenkins, Henry (1992). *Textual Poachers: Television Fans and Participatory Culture*. Oxford, UNITED KINGDOM: Routledge. ISBN: 978-1-136-29072-5. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=1097854 (visited on 06/03/2024).

Jensen, Jakob Linaa and Sander Andreas Schwartz (Oct. 2021). "The Return of the "Lurker": A Longitudinal Study of Citizens' Use of Social Media in Danish Elections 2011, 2015, and 2019". en. In: *Social Media + Society* 7.4, p. 20563051211063463. ISSN: 2056-3051. DOI: 10.1177/20563051211063463. URL: https://doi.org/10.1177/20563051211063463 (visited on 05/18/2023).

Jeong, H., Z. Néda, and A. L. Barabási (Feb. 2003). "Measuring preferential attachment in evolving networks". en. In: *Europhysics Letters* 61.4, p. 567. ISSN: 0295-5075. DOI: 10.1209/epl/i2003-00166-9. URL: https://iopscience.iop.org/article/10.1209/epl/i2003-00166-9/meta (visited on 04/24/2024).

Jingnan, Huo (Feb. 2023). "Twitter's new data access rules will make social media research harder". en. In: *NPR*. URL: https://www.npr.org/2023/02/09/1155543369/twitters-new-data-access-rules-will-make-social-media-research-harder (visited on 08/02/2023).

Jüni, Peter et al. (Aug. 2021). "Revisiting the evidence for physical distancing, face masks, and eye protection". English. In: *The Lancet* 398.10301, p. 663. ISSN: 0140-6736, 1474-547X. DOI: 10.1016/

S0140-6736(21)01758-X. URL: https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(21)01758-X/fulltext (visited on 12/28/2021).

Just, Sine N. (Sept. 2019). "An assemblage of avatars: Digital organization as affective intensification in the GamerGate controversy". en. In: *Organization* 26.5, pp. 716–738. ISSN: 1350-5084. DOI: 10.1177/1350508419842710. URL: https://doi.org/10.1177/1350508419842710 (visited on 05/02/2024).

Kahan, Dan M., Hank Jenkins-Smith, and Donald Braman (Feb. 2011). "Cultural cognition of scientific consensus". In: *Journal of Risk Research* 14.2, pp. 147–174. ISSN: 1366-9877. DOI: 10.1080/13669877.2010.511246. URL: https://doi.org/10.1080/13669877.2010.511246 (visited on 06/18/2021).

Kahle, Brewster (1997). *Preserving the internet*. Vol. 276.

Karell, Daniel et al. (Apr. 2023). ""Born for a Storm": Hard-Right Social Media and Civil Unrest". en. In: *American Sociological Review* 88.2, pp. 322–349. ISSN: 0003-1224. DOI: 10.1177/00031224231156190. URL: https://doi.org/10.1177/00031224231156190 (visited on 05/07/2024).

Katz, Elihu and Paul F. Lazarsfeld (1955). *Personal influence: the part played by people in the flow of mass communications*. Personal influence: the part played by people in the flow of mass communications. New York, NY, US: Free Press.

Kawintiranon, Kornraphop and Lisa Singh (2022). "PoliBERTweet: a pre-trained language model for analyzing political content on twitter". In: *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pp. 7360–7367. URL: https://aclanthology.org/2022.lrec-1.801/ (visited on 02/27/2024).

Kennedy, Ian et al. (June 2022). "Repeat Spreaders and Election Delegitimization: A Comprehensive Dataset of Misinformation Tweets from the 2020 U.S. Election". en. In: *Journal of Quantitative Description: Digital Media* 2. ISSN: 2673-8813. DOI: 10.51685/jqd.2022.013. URL: https://journalqd.org/article/view/3137 (visited on 12/08/2022).

Kiene, Charles, Andrés Monroy-Hernández, and Benjamin Mako Hill (2016). "Surviving an" eternal september" how an online community managed a surge of newcomers". In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 1152–1156.

Klassen, Shamika and Casey Fiesler (Oct. 2022). ""This Isn't Your Data, Friend": Black Twitter as a Case Study on Research Ethics for Public Data". en. In: *Social Media + Society* 8.4, p. 205630512211443. ISSN: 2056-3051, 2056-3051. DOI: 10.1177/20563051221144317. URL: http://journals.sagepub.com/doi/10.1177/20563051221144317 (visited on 04/21/2024).

Kleinberg, Jon M (1999). "Hubs, authorities, and communities". In: *ACM computing surveys (CSUR)* 31.4es, 5–es.

Klompas, Michael et al. (May 2020). "Universal Masking in Hospitals in the Covid-19 Era". In: *New England Journal of Medicine* 382.21, e63. ISSN: 0028-4793. DOI: 10.1056/nejmp2006372. URL: https://doi.org/10.1056\%2Fnejmp2006372 (visited on 06/10/2023).

Klonick, Kate (2015). "Re-Shaming the Debate: Social Norms, Shame, and Regulation in an Internet Age". In: *Maryland Law Review* 75, p. 1029. URL: https://heinonline.org/HOL/Page?handle=hein.journals/mllr75\&id=1053\&div=\&collection=.

Kofman, Ava (Oct. 2018). "Bruno Latour, the Post-Truth Philosopher, Mounts a Defense of Science". en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/

2018/10/25/magazine/bruno-latour-post-truth-philosopher-science.html (visited on 08/29/2022).

Koltai, Kolina S. and Kenneth R. Fleischmann (2017). "Questioning science with science: The evolution of the vaccine safety movement". en. In: *Proceedings of the Association for Information Science and Technology* 54.1, pp. 232–240. ISSN: 2373-9231. DOI: 10.1002/pra2.2017.14505401026. URL: http://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/pra2.2017.14505401026 (visited on 11/02/2020).

Kossinets, Gueorgi and Duncan J. Watts (Sept. 2009). "Origins of Homophily in an Evolving Social Network". In: *American Journal of Sociology* 115.2, pp. 405–450. ISSN: 0002-9602. DOI: 10.1086/599247. URL: https://www.journals.uchicago.edu/doi/full/10.1086/599247 (visited on 05/08/2023).

Kraut, Robert E. and Paul Resnick (2012). *Building successful online communities: Evidence-based social design*. Mit Press.

Kreiss, Daniel and Shannon C McGregor (Jan. 2024). "A review and provocation: On polarization and platforms". en. In: *New Media & Society* 26.1, pp. 556–579. ISSN: 1461-4448. DOI: 10.1177/14614448231161880. URL: https://doi.org/10.1177/14614448231161880 (visited on 04/16/2024).

Kunegis, Jérôme, Andreas Lommatzsch, and Christian Bauckhage (2009). "The Slashdot Zoo: Mining a Social Network with Negative Edges". In: *Proceedings of the 18th international conference on World wide web - WWW '09*, p. 741. DOI: 10.1145/1526709.1526809. URL: http://arxiv.org/abs/1710.11395 (visited on 12/18/2022).

Larsson, Anders Olof and Hallvard Moe (2015). "From Emerging to Established?: A Comparison of Twitter Use during Swedish Election Campaigns in 2010 and 2014". In: *The Routledge companion to social media and politics*. Routledge, pp. 311–324.

Lawson, Caitlin E. (June 2018). "Platform vulnerabilities: harassment and misogynoir in the digital attack on Leslie Jones". In: *Information, Communication & Society* 21.6, pp. 818–833. ISSN: 1369-118X. DOI: 10.1080/1369118X.2018.1437203. URL: https://doi.org/10.1080/1369118X.2018.1437203 (visited on 03/24/2023).

LeBlanc, Paul and Clare Duffy (Nov. 2022). "Elon Musk restores Donald Trump's Twitter account". en. In: *CNN*. URL: https://www.cnn.com/2022/11/19/business/twitter-musk-trump-reinstate/index.html (visited on 04/08/2024).

Lee, Crystal et al. (May 2021). "Viral Visualizations: How Coronavirus Skeptics Use Orthodox Data Practices to Promote Unorthodox Science Online". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* 607, pp. 1–18. URL: https://doi.org/10.1145/3411764.3445211 (visited on 10/27/2021).

Lee, Nicole M. and Matthew S. VanDyke (Aug. 2015). "Set It and Forget It: The One-Way Use of Social Media by Government Agencies Communicating Science". en. In: *Science Communication* 37.4, pp. 533–541. ISSN: 1075-5470. DOI: 10.1177/1075547015588600. URL: https://doi.org/10.1177/1075547015588600 (visited on 08/26/2022).

Leffler, Christopher T. et al. (Oct. 2020). "Association of Country-wide Coronavirus Mortality with Demographics, Testing, Lockdowns, and Public Wearing of Masks". EN. In: *The American Journal of Tropical Medicine and Hygiene* 103.6, pp. 2400–2411. ISSN: 0002-9637, 1476-1645. DOI: 10.4269/

ajtmh.20-1015. URL: https://www.ajtmh.org/view/journals/tpmd/103/6/article-p2400.xml (visited on 06/15/2022).

Leingang, Rachel (Jan. 2024). "'Stakes are really high': misinformation researcher changes tack for 2024 US election". en-GB. In: *The Guardian*. ISSN: 0261-3077. URL: https://www.theguardian.com/us-news/2024/jan/01/misinformation-trends-2024-election-right-wing (visited on 04/30/2024).

Leskovec, Jure, Daniel Huttenlocher, and Jon Kleinberg (Apr. 2010). "Signed networks in social media". en. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Atlanta Georgia USA: ACM, pp. 1361–1370. ISBN: 978-1-60558-929-9. DOI: 10.1145/1753326.1753532. URL: https://dl.acm.org/doi/10.1145/1753326.1753532 (visited on 04/13/2024).

Leung, Nancy H. L. et al. (May 2020). "Respiratory virus shedding in exhaled breath and efficacy of face masks". en. In: *Nature Medicine* 26.5, pp. 676–680. ISSN: 1546-170X. DOI: 10.1038/s41591-020-0843-2. URL: https://www.nature.com/articles/s41591-020-0843-2 (visited on 12/02/2021).

Lewandowsky, Stephan, Gilles E. Gignac, and Samuel Vaughan (Apr. 2013). "The pivotal role of perceived scientific consensus in acceptance of science". en. In: *Nature Climate Change* 3.4, pp. 399–404. ISSN: 1758-6798. DOI: 10.1038/nclimate1720. URL: https://www.nature.com/articles/nclimate1720 (visited on 08/29/2022).

Lewandowsky, Stephan and Sander van der Linden (July 2021). "Countering Misinformation and Fake News Through Inoculation and Prebunking". In: *European Review of Social Psychology* 32.2, pp. 348–384. ISSN: 1046-3283. DOI: 10.1080/10463283.2021.1876983. URL: https://doi.org/10.1080/10463283.2021.1876983 (visited on 08/29/2022).

Lewandowsky, Stephan et al. (Nov. 2019). "Science by social media: Attitudes towards climate change are mediated by perceived social consensus". en. In: *Memory & Cognition* 47.8, pp. 1445–1456. ISSN: 1532-5946. DOI: 10.3758/s13421-019-00948-y. URL: https://doi.org/10.3758/s13421-019-00948-y (visited on 08/26/2022).

Lewis, Rebecca (Sept. 2018). *Alternative influence: Broadcasting the reactionary right on YouTube*. Tech. rep. Data & Society Research Institute. URL: https://apo.org.au/sites/default/files/resource-files/2018-09/apo-nid193281.pdf (visited on 04/16/2024).

– (Feb. 2020). ""This Is What the News Won't Show You": YouTube Creators and the Reactionary Politics of Micro-celebrity". en. In: *Television & New Media* 21.2, pp. 201–217. ISSN: 1527-4764. DOI: 10.1177/1527476419879919. URL: https://doi.org/10.1177/1527476419879919 (visited on 12/09/2022).

Lewis, Rebecca and Angèle Christin (July 2022). "Platform drama: "Cancel culture," celebrity, and the struggle for accountability on YouTube". en. In: *New Media & Society* 24.7, pp. 1632–1656. ISSN: 1461-4448, 1461-7315. DOI: 10.1177/14614448221099235. URL: http://journals.sagepub.com/doi/10.1177/14614448221099235 (visited on 04/15/2024).

Li, E. Rosalie and Alex Kaplan (Dec. 2021). *Twitter has allowed one user to manipulate the platform by coordinating over 40 hashtags, seemingly in violation of policy*. en. URL: https://www.mediamatters.org/twitter/twitter-has-allowed-one-user-manipulate-platform-coordinating-over-40-hashtags-seemingly (visited on 03/03/2024).

Liang, Calvin A., Sean A. Munson, and Julie A. Kientz (Apr. 2021). "Embracing Four Tensions in Human-Computer Interaction Research with Marginalized People". en. In: *ACM Transactions on*

Computer-Human Interaction 28.2, pp. 1–47. ISSN: 1073-0516, 1557-7325. DOI: 10.1145/3443686. URL: https://dl.acm.org/doi/10.1145/3443686 (visited on 04/12/2024).

Lilleker, Darren, Karolina Koc-Michalska, and Nigel Jackson (2015). "Social media in the UK election campaigns 2008-14: Experimentation, innovation and convergence". In.

Linden, Sander van der, Jon Roozenbeek, and Josh Compton (2020). "Inoculating Against Fake News About COVID-19". In: *Frontiers in Psychology* 11, p. 2928. ISSN: 1664-1078. DOI: 10.3389/fpsyg.2020.566790. URL: https://www.frontiersin.org/article/10.3389/fpsyg.2020.566790 (visited on 12/31/2021).

Lippmann, Walter (1922). *Public opinion*. Public opinion. New York, NY, US: MacMillan Co. DOI: 10.1037/14847-000.

Litt, Eden (July 2012). "Knock, Knock. Who's There? The Imagined Audience". In: *Journal of Broadcasting & Electronic Media* 56.3, pp. 330–345. ISSN: 0883-8151. DOI: 10.1080/08838151.2012.705195. URL: https://doi.org/10.1080/08838151.2012.705195 (visited on 05/29/2024).

Liu, Amy (Oct. 2011). "Unraveling the myth of meritocracy within the context of US higher education". en. In: *Higher Education* 62.4, pp. 383–397. ISSN: 1573-174X. DOI: 10.1007/s10734-010-9394-7. URL: https://doi.org/10.1007/s10734-010-9394-7 (visited on 05/04/2024).

Livingstone, Sonia (2005). "On the relation between audiences and publics". In: *Audiences and Publics: When Cultural Engagement Matters for the Public Sphere*. Bristol, UNITED KINGDOM: Intellect, Limited. ISBN: 978-1-84150-923-5. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=283034 (visited on 05/29/2024).

Long, Youlin et al. (May 2020). "Effectiveness of N95 respirators versus surgical masks against influenza: A systematic review and meta-analysis". eng. In: *Journal of Evidence-Based Medicine* 13.2, pp. 93–101. ISSN: 1756-5391. DOI: 10.1111/jebm.12381.

Lorde, Audre (1997). "The uses of anger". In: *Women's Studies Quarterly* 25.1/2, pp. 278–285.

Lorenz, Taylor (Dec. 2022). "Meet the woman behind Libs of TikTok, secretly fueling the right's outrage machine". en-US. In: *Washington Post*. ISSN: 0190-8286. URL: https://www.washingtonpost.com/technology/2022/04/19/libs-of-tiktok-right-wing-media/ (visited on 04/11/2024).

Lose your Fear of Academic Twitter (Mar. 2021). Panel Discussion. URL: https://www.youtube.com/watch?v=2dL1ZfDy9kQ (visited on 05/05/2024).

Lu, Jessica H. and Catherine Knight Steele (May 2019). "'Joy is resistance': cross-platform resilience and (re)invention of Black oral culture online". In: *Information, Communication & Society* 22.6, pp. 823–837. ISSN: 1369-118X. DOI: 10.1080/1369118X.2019.1575449. URL: https://doi.org/10.1080/1369118X.2019.1575449 (visited on 05/04/2024).

Lukito, Josephine (Mar. 2020). "Coordinating a Multi-Platform Disinformation Campaign: Internet Research Agency Activity on Three U.S. Social Media Platforms, 2015 to 2017". In: *Political Communication* 37.2, pp. 238–255. DOI: 10.1080/10584609.2019.1661889. (Visited on 01/16/2022).

Lyu, Wei and George L. Wehby (Aug. 2020). "Community Use Of Face Masks And COVID-19: Evidence From A Natural Experiment Of State Mandates In The US". In: *Health Affairs* 39.8, pp. 1419–1425. ISSN: 0278-2715. DOI: 10.1377/hlthaff.2020.00818. URL: https://www.healthaffairs.org/doi/10.1377/hlthaff.2020.00818 (visited on 12/30/2021).

Ma, Cindy (July 2021). "What is the "lite" in "alt-lite?" The discourse of white vulnerability and dominance among YouTube's reactionaries". en. In: *Social Media + Society* 7.3, p. 20563051211036385.

ISSN: 2056-3051. DOI: 10.1177/20563051211036385. URL: https://doi.org/10.1177/20563051211036385 (visited on 05/05/2024).

Ma, Renkai, Xinning Gui, and Yubo Kou (2023). "Multi-Platform Content Creation: The Configuration of Creator Ecology through Platform Prioritization, Content Synchronization, and Audience Management". In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–19.

Ma, Renkai and Yubo Kou (Nov. 2022). ""I am not a YouTuber who can make whatever video I want. I have to keep appeasing algorithms": Bureaucracy of Creator Moderation on YouTube". In: *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social Computing*. CSCW'22 Companion. New York, NY, USA: Association for Computing Machinery, pp. 8–13. ISBN: 978-1-4503-9190-0. DOI: 10.1145/3500868.3559445. URL: https://dl.acm.org/doi/10.1145/3500868.3559445 (visited on 04/20/2024).

MacIntyre, C. Raina and Nancy Binkin (Jan. 2021). "In the room where it happens: The consequences of the lack of public health expertise during the COVID-19 pandemic". en. In: *Global Biosecurity* 3.1. ISSN: 2652-0036. DOI: 10.31646/gbio.102. URL: http://jglobalbiosecurity.com/articles/10.31646/gbio.102/ (visited on 12/27/2021).

MacIntyre, C Raina and Abrar A. Chughtai (Nov. 2020a). "Masks in the community are an effective strategy: Author's response to Haslam et al (2020)". In: *International Journal of Nursing Studies* 111, p. 103751. ISSN: 0020-7489. DOI: 10.1016/j.ijnurstu.2020.103751. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7480396/ (visited on 12/27/2021).

MacIntyre, C. Raina and Abrar Ahmad Chughtai (Aug. 2020b). "A rapid systematic review of the efficacy of face masks and respirators against coronaviruses and other respiratory transmissible viruses for the community, healthcare workers and sick patients". en. In: *International Journal of Nursing Studies* 108, p. 103629. ISSN: 0020-7489. DOI: 10.1016/j.ijnurstu.2020.103629. URL: https://www.sciencedirect.com/science/article/pii/S0020748920301139 (visited on 12/27/2021).

MacIntyre, C Raina et al. (Apr. 2015). "A cluster randomised trial of cloth masks compared with medical masks in healthcare workers". In: *BMJ Open* 5.4, e006577. ISSN: 2044-6055. DOI: 10.1136/bmjopen-2014-006577. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4420971/ (visited on 10/25/2022).

MacIntyre, C. Raina et al. (Nov. 2022). "Effectiveness of facemasks for opening a university campus in Mississippi, United States – a modelling study". In: *Journal of American College Health* 70.8, pp. 2505–2510. ISSN: 0744-8481. DOI: 10.1080/07448481.2020.1866579. URL: https://doi.org/10.1080/07448481.2020.1866579 (visited on 05/25/2023).

MacIntyre, Chandini Raina et al. (Sept. 2020). "Contamination and washing of cloth masks and risk of infection among hospital health workers in Vietnam: a post hoc analysis of a randomised controlled trial". en. In: *BMJ Open* 10.9, e042045. ISSN: 2044-6055, 2044-6055. DOI: 10.1136/bmjopen-2020-042045. URL: https://bmjopen.bmj.com/content/10/9/e042045 (visited on 12/27/2021).

MacIntyre, Chandini Raina et al. (May 2021). "Mask use, risk-mitigation behaviours and pandemic fatigue during the COVID-19 pandemic in five cities in Australia, the UK and USA: A cross-sectional survey". en. In: *International Journal of Infectious Diseases* 106, pp. 199–207. ISSN: 1201-9712. DOI: 10.1016/j.ijid.2021.03.056. URL: https://www.sciencedirect.com/science/article/pii/S1201971221002745 (visited on 12/26/2021).

Maddock, Jim et al. (Feb. 2015). "Characterizing Online Rumoring Behavior Using Multi-Dimensional Signatures". en. In: *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. Vancouver BC Canada: ACM, pp. 228–241. ISBN: 978-1-4503-2922-4. DOI: 10.1145/2675133.2675280. URL: https://dl.acm.org/doi/10.1145/2675133.2675280 (visited on 05/01/2024).

Majó-Vázquez, Silvia et al. (2021). "The role of suspended accounts in political discussion on social media: Analysis of the 2017 French, UK and German elections". In: *Social Media+ Society* 7.3, p. 20563051211027202.

Malinen, Sanna (2015). "Understanding user participation in online communities: A systematic literature review of empirical studies". In: *Computers in human behavior* 46, pp. 228–238.

Marres, Noortje (2015). "Why map issues? On controversy analysis as a digital method". In: *Science, Technology, & Human Values* 40.5, pp. 655–686.

Marwick, Alice (2013a). "They're really profound women, they're entrepreneurs': Conceptions of authenticity in fashion blogging". In: *7th international AIII conference on weblogs and social media (ICWSM)*. Vol. 2011, pp. 1–8. URL: https://www.tiara.org/wp-content/uploads/2018/05/amarwick_fashionblogs_ICWSM_2013.pdf (visited on 05/07/2024).

Marwick, Alice and Danah boyd (May 2011). "To See and Be Seen: Celebrity Practice on Twitter". en. In: *Convergence: The International Journal of Research into New Media Technologies* 17.2, pp. 139–158. ISSN: 1354-8565, 1748-7382. DOI: 10.1177/1354856510394539. URL: http://journals.sagepub.com/doi/10.1177/1354856510394539 (visited on 11/07/2023).

Marwick, Alice E. (2010). "Status update: Celebrity, publicity and self-branding in Web 2.0". English. Ph.D. United States – New York: New York University. URL: https://www.proquest.com/docview/763612310/abstract/CE8A441D198747ECPQ/1 (visited on 04/18/2024).

– (2013b). *Status Update: Celebrity, Publicity, and Branding in the Social Media Age*. New Haven, UNITED STATES: Yale University Press. ISBN: 978-0-300-19915-4. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=3421330 (visited on 05/07/2024).

– (Apr. 2021). "Morally Motivated Networked Harassment as Normative Reinforcement". en. In: *Social Media + Society* 7.2, p. 20563051211021378. ISSN: 2056-3051. DOI: 10.1177/20563051211021378. URL: https://doi.org/10.1177/20563051211021378 (visited on 03/24/2023).

Marwick, Alice E. and Danah Boyd (Feb. 2011). "I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience". en. In: *New Media & Society* 13.1, pp. 114–133. ISSN: 1461-4448, 1461-7315. DOI: 10.1177/1461444810365313. URL: http://journals.sagepub.com/doi/10.1177/1461444810365313 (visited on 06/03/2024).

Mason, Shannon and Lenandlar Singh (Apr. 2022). "Reporting and discoverability of "Tweets" quoted in published scholarship: current practice and ethical implications". en. In: *Research Ethics* 18.2, pp. 93–113. ISSN: 1747-0161. DOI: 10.1177/17470161221076948. URL: https://doi.org/10.1177/17470161221076948 (visited on 04/21/2024).

Massanari, Adrienne (Mar. 2017). "#Gamergate and The Fappening: How Reddit's algorithm, governance, and culture support toxic technocultures". en. In: *New Media & Society* 19.3, pp. 329–346. ISSN: 1461-4448. DOI: 10.1177/1461444815608807. URL: https://doi.org/10.1177/1461444815608807 (visited on 03/24/2023).

Mastrangelo, Frankie and Gina Marie Longo (Jan. 2024). "Downlining Disinformation: How MLM Distributors Use Gendered Strategies for Recruitment and Pastel QAnon Indoctrination". en. In: *Social Media + Society* 10.1, p. 20563051231224735. ISSN: 2056-3051. DOI: 10.1177/20563051231224735. URL: https://doi.org/10.1177/20563051231224735 (visited on 04/20/2024).

Matalon, Yogev et al. (Mar. 2021). "Using sentiment analysis to predict opinion inversion in Tweets of political communication". en. In: *Scientific Reports* 11.1, p. 7250. ISSN: 2045-2322. DOI: 10.1038/s41598-021-86510-w. URL: https://www.nature.com/articles/s41598-021-86510-w (visited on 04/13/2024).

Matzko, Paul (2020). *The radio right: How a band of broadcasters took on the federal government and built the modern conservative movement*. Oxford University Press, USA.

McCulloh, Ian and Kathleen M. Carley (2011). "Detecting change in longitudinal social networks". In: *Journal of social structure* 12.1, pp. 1–37.

McGregor, Shannon C (Aug. 2019). "Social media as public opinion: How journalists use social media to represent public opinion". en. In: *Journalism* 20.8, pp. 1070–1086. ISSN: 1464-8849. DOI: 10.1177/1464884919845458. URL: https://doi.org/10.1177/1464884919845458 (visited on 07/30/2023).

McKelvey, Karissa, Joseph DiGrazia, and Fabio Rojas (Apr. 2014). "Twitter publics: how online political communities signaled electoral outcomes in the 2010 US house election". In: *Information, Communication & Society* 17.4, pp. 436–450. ISSN: 1369-118X. DOI: 10.1080/1369118X.2014.892149. URL: https://doi.org/10.1080/1369118X.2014.892149 (visited on 05/16/2023).

McNair, Brian (2017). *An introduction to political communication*. Routledge. URL: https://www.taylorfrancis.com/books/mono/10.4324/9781315750293/introduction-political-communication-brian-mcnair-brian-mcnair (visited on 04/17/2024).

McNear, Claire (May 2018). "How Quote Tweets Helped Ruin Twitter". en. In: *The Ringer*. URL: https://www.theringer.com/tech/2018/5/2/17311616/twitter-retweet-quote-endorsement-function-trolls (visited on 04/11/2024).

McRae, Elizabeth Gillespie (2018). *Mothers of Massive Resistance: White Women and the Politics of White Supremacy*. en. Oxford University Press. ISBN: 978-0-19-027171-8.

McRae, Sarah (Nov. 2020). "Get off my internets: How anti-fans deconstruct lifestyle bloggers' authenticity work". In: *Persona Studies* 3.1, pp. 13–27. DOI: 10.3316/informit.956534970080756. URL: https://search.informit.org/doi/abs/10.3316/informit.956534970080756 (visited on 06/03/2024).

Megill, Allan (June 1989). "Recounting the Past: "Description," Explanation, and Narrative in Historiography". In: *The American Historical Review* 94.3, pp. 627–653. ISSN: 0002-8762. DOI: 10.1086/ahr/94.3.627. URL: https://doi.org/10.1086/ahr/94.3.627 (visited on 04/29/2024).

Meisner, Colten (Dec. 2023). "The weaponization of platform governance: Mass reporting and algorithmic punishments in the creator economy". en. In: *Policy & Internet* 15.4, pp. 466–477. ISSN: 1944-2866, 1944-2866. DOI: 10.1002/poi3.359. URL: https://onlinelibrary.wiley.com/doi/10.1002/poi3.359 (visited on 04/21/2024).

Mejia, Robert, Kay Beckermann, and Curtis Sullivan (Apr. 2018). "White lies: a racial history of the (post)truth". In: *Communication and Critical/Cultural Studies* 15.2, pp. 109–126. DOI: 10.1080/14791420.2018.1456668. (Visited on 01/16/2022).

Mellado, Claudia and Alfred Hermida (Jan. 2021). "The Promoter, Celebrity, and Joker Roles in Journalists' Social Media Performance". en. In: *Social Media + Society* 7.1, p. 2056305121990643. ISSN: 2056-3051. DOI: 10.1177/2056305121990643. URL: https://doi.org/10.1177/2056305121990643 (visited on 04/20/2024).

Messeri, Lisa and M. J. Crockett (Mar. 2024). "Artificial intelligence and illusions of understanding in scientific research". en. In: *Nature* 627.8002, pp. 49–58. ISSN: 1476-4687. DOI: 10.1038/s41586-024-07146-0. URL: https://www.nature.com/articles/s41586-024-07146-0 (visited on 05/04/2024).

Metaxas, Panagiotis et al. (2015). "What do retweets indicate? Results from user survey and meta-review of research". In: *Proceedings of the international AAAI conference on web and social media*. Vol. 9, pp. 658–661.

Michaeli, Ethan (2016). *The Defender: How the Legendary Black Newspaper Changed America*. Houghton Mifflin Harcourt. (Visited on 11/06/2023).

Miller, Boaz (2019). "The Social Epistemology of Consensus and Dissent". In: *The Routledge Handbook of Social Epistemology*. Routledge. ISBN: 978-1-315-71793-7.

Miller, Kaitlin C. (Aug. 2023). "Hostility Toward the Press: A Synthesis of Terms, Research, and Future Directions in Examining Harassment of Journalists". In: *Digital Journalism* 11.7, pp. 1230–1249. ISSN: 2167-0811. DOI: 10.1080/21670811.2021.1991824. URL: https://doi.org/10.1080/21670811.2021.1991824 (visited on 04/12/2024).

Mingers, John and Loet Leydesdorff (Oct. 2015). "A review of theory and practice in scientometrics". en. In: *European Journal of Operational Research* 246.1, pp. 1–19. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2015.04.002. URL: https://www.sciencedirect.com/science/article/pii/S037722171500274X (visited on 08/29/2022).

Mohammad, Saif et al. (2016). "Semeval-2016 task 6: Detecting stance in tweets". In: *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)*, pp. 31–41.

Moran, Rachel E., Izzi Grasso, and Kolina Koltai (Oct. 2022). "Folk Theories of Avoiding Content Moderation: How Vaccine-Opposed Influencers Amplify Vaccine Opposition on Instagram". en. In: *Social Media + Society* 8.4, p. 20563051221144252. ISSN: 2056-3051. DOI: 10.1177/20563051221144252. URL: https://doi.org/10.1177/20563051221144252 (visited on 05/10/2023).

Moran, Rachel E., Anna L. Swan, and Taylor Agajanian (2024). "Vaccine Misinformation for Profit: Conspiratorial Wellness Influencers and the Monetization of Alternative Health". In: *International Journal of Communication* 18, p. 23. URL: https://ijoc.org/index.php/ijoc/article/view/21128 (visited on 04/20/2024).

Moravcsik, Michael J. and Poovanalingam Murugesan (1975). "Some Results on the Function and Quality of Citations". In: *Social Studies of Science* 5.1, pp. 86–92. ISSN: 0306-3127. URL: https://www.jstor.org/stable/284557 (visited on 05/26/2022).

Morgan, Allison C. et al. (Dec. 2022). "Socioeconomic roots of academic faculty". en. In: *Nature Human Behaviour* 6.12, pp. 1625–1633. ISSN: 2397-3374. DOI: 10.1038/s41562-022-01425-4. URL: https://www.nature.com/articles/s41562-022-01425-4 (visited on 05/04/2024).

Morley, David (Dec. 1993). "Active Audience Theory: Pendulums and Pitfalls". In: *Journal of Communication* 43.4, pp. 13–19. ISSN: 0021-9916. DOI: 10.1111/j.1460-2466.1993.tb01299.x. URL: https://doi.org/10.1111/j.1460-2466.1993.tb01299.x (visited on 05/29/2024).

Mortensen, Torill Elvira (Dec. 2018). "Anger, Fear, and Games: The Long Event of #GamerGate". en. In: *Games and Culture* 13.8, pp. 787–806. ISSN: 1555-4120, 1555-4139. DOI: 10.1177/1555412016640408. URL: http://journals.sagepub.com/doi/10.1177/1555412016640408 (visited on 04/12/2024).

Muir, Shannon R., Lynne D. Roberts, and Lorraine P. Sheridan (Jan. 2021). "The portrayal of online shaming in contemporary online news media: A media framing analysis". In: *Computers in Human Behavior Reports* 3, p. 100051. ISSN: 2451-9588. DOI: 10.1016/j.chbr.2020.100051. URL: https://www.sciencedirect.com/science/article/pii/S2451958820300518 (visited on 04/16/2024).

Munger, Kevin and Joseph Phillips (Jan. 2022). "Right-Wing YouTube: A Supply and Demand Perspective". en. In: *The International Journal of Press/Politics* 27.1, pp. 186–219. ISSN: 1940-1612. DOI: 10.1177/1940161220964767. URL: https://doi.org/10.1177/1940161220964767 (visited on 05/01/2024).

Musgrave, Tyler, Alia Cummings, and Sarita Schoenebeck (Apr. 2022). "Experiences of Harm, Healing, and Joy among Black Women and Femmes on Social Media". In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New York, NY, USA: Association for Computing Machinery, pp. 1–17. ISBN: 978-1-4503-9157-3. DOI: 10.1145/3491102.3517608. URL: https://dl.acm.org/doi/10.1145/3491102.3517608 (visited on 05/02/2024).

Nakamura, Lisa (2015). "The unwanted labour of social media: Women of colour call out culture as venture community management". In: *New Formations* 86.86, pp. 106–112.

Napoli, Philip M. (2010). "Contextualizing Audience Evolution". In: *Audience Evolution: New Technologies and the Transformation of Media Audiences*. New York, UNITED STATES: Columbia University Press. ISBN: 978-0-231-52094-2. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=895121 (visited on 06/03/2024).

Nelson, Jacob L. and James G. Webster (Jan. 2016). "Audience Currencies in the Age of Big Data". en. In: *International Journal on Media Management* 18.1, pp. 9–24. ISSN: 1424-1277, 1424-1250. DOI: 10.1080/14241277.2016.1166430. URL: http://www.tandfonline.com/doi/full/10.1080/14241277.2016.1166430 (visited on 06/03/2024).

Newell, Edward et al. (2016). "User Migration in Online Social Networks: A Case Study on Reddit During a Period of Community Unrest". en. In: *Proceedings of the International AAAI Conference on Web and Social Media* 10.1, pp. 279–288. ISSN: 2334-0770. DOI: 10.1609/icwsm.v10i1.14750. URL: https://ojs.aaai.org/index.php/ICWSM/article/view/14750 (visited on 05/09/2023).

Newman, M. E. J. (June 2001). "Scientific collaboration networks. II. Shortest paths, weighted networks, and centrality". In: *Physical Review E* 64.1, p. 016132. DOI: 10.1103/PhysRevE.64.016132. (Visited on 08/03/2021).

– (June 2006). "Modularity and community structure in networks". In: *Proceedings of the National Academy of Sciences* 103.23, pp. 8577–8582. DOI: 10.1073/pnas.0601602103. URL: https://www.pnas.org/doi/abs/10.1073/pnas.0601602103 (visited on 04/24/2024).

Newman, M. E. J. and M. Girvan (Feb. 2004). "Finding and evaluating community structure in networks". In: *Physical Review E* 69.2, p. 026113. DOI: 10.1103/PhysRevE.69.026113. URL: https://link.aps.org/doi/10.1103/PhysRevE.69.026113 (visited on 12/16/2021).

Newman, Nic et al. (June 2023). *Digital News Report 2023*. en. Tech. rep. Reuters Institute for the Study of Journalism. URL: https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2023 (visited on 11/06/2023).

Nguyễn, Sarah et al. (Mar. 2022). "Studying mis- and disinformation in Asian diasporic communities: The need for critical transnational research beyond Anglocentrism". en-US. In: *Harvard Kennedy School Misinformation Review*. DOI: 10.37016/mr-2020-95. URL: https://misinforeview.hks.harvard.edu/article/studying-mis-and-disinformation-in-asian-diasporic-communities-the-need-for-critical-transnational-research-beyond-anglocentrism/ (visited on 04/20/2024).

Nguyen, Dat Quoc, Thanh Vu, and Anh Tuan Nguyen (Oct. 2020). "BERTweet: A pre-trained language model for English Tweets". In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Ed. by Qun Liu and David Schlangen. Online: Association for Computational Linguistics, pp. 9–14. DOI: 10.18653/v1/2020.emnlp-demos.2. URL: https://aclanthology.org/2020.emnlp-demos.2 (visited on 03/19/2024).

Nick, Bobo et al. (Aug. 2013). "Simmelian backbones: amplifying hidden homophily in Facebook networks". In: *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. ASONAM '13. New York, NY, USA: Association for Computing Machinery, pp. 525–532. ISBN: 978-1-4503-2240-9. DOI: 10.1145/2492517.2492569. URL: https://dl.acm.org/doi/10.1145/2492517.2492569 (visited on 05/08/2024).

Nieborg, David B, Brooke Erin Duffy, and Thomas Poell (July 2020). "Studying Platforms and Cultural Production: Methods, Institutions, and Practices". en. In: *Social Media + Society* 6.3, p. 2056305120943273. ISSN: 2056-3051. DOI: 10.1177/2056305120943273. URL: https://doi.org/10.1177/2056305120943273 (visited on 05/05/2024).

Nocaj, Arlind, Mark Ortmann, and Ulrik Brandes (2015). "Untangling the Hairballs of Multi-Centered, Small-World Online Social Media Networks". en. In: *Journal of Graph Algorithms and Applications* 19.2, pp. 595–618. ISSN: 1526-1719. DOI: 10.7155/jgaa.00370. URL: http://jgaa.info/getPaper?id=370 (visited on 06/17/2021).

nuanceORDEATH (July 2020). *#MASKUP THREAD*. en. Tweet. URL: https://twitter.com/nuanceORDEATH/status/1279144399897866248 (visited on 12/01/2022).

Obar, Jonathan A. and Anne Oeldorf-Hirsch (June 2018). *The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services*. en. SSRN Scholarly Paper ID 2757465. Rochester, NY: Social Science Research Network. DOI: 10.2139/ssrn.2757465. URL: https://papers.ssrn.com/abstract=2757465 (visited on 06/11/2021).

Oliver, Pamela (1980). "Rewards and Punishments as Selective Incentives for Collective Action: Theoretical Investigations". In: *American Journal of Sociology* 85.6, pp. 1356–1375. ISSN: 0002-9602. URL: https://www.jstor.org/stable/2778382 (visited on 04/16/2024).

Oreskes, Naomi (2021). *Why trust science?* Princeton University Press.

Oreskes, Naomi and Erik M Conway (2012). *Merchants of doubt: how a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. English. London: Bloomsbury. ISBN: 978-1-4088-2483-2.

Ovide, Shira (Jan. 2021). "Trump Isn't the Only One". en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2021/01/08/technology/trump-misinformation-superspreaders.html (visited on 09/19/2023).

Papacharissi, Zizi (2015). *Affective publics: Sentiment, technology, and politics*. Oxford University Press.

Partisan Conflict and Congressional Outreach (Feb. 2017). en-US. Tech. rep. Pew Research Center. URL: https://www.pewresearch.org/politics/2017/02/23/partisan-conflict-and-congressional-outreach/ (visited on 04/16/2024).

Pasquetto, Irene V et al. (2022). "Disinformation as Infrastructure: Making and Maintaining the QAnon Conspiracy on Italian Digital Media". In: *Proceedings of the ACM on Human-Computer Interaction* 6.CSCW1, pp. 1–31.

Pavlenko, Aneta (Dec. 2018). "Superdiversity and Why It Isn't: Reflections on Terminological Innovation and Academic Branding". In: *Sloganization in Language Education Discourse*. Ed. by Barbara Schmenk, Stephan Breidbach, and Lutz Küster. Bristol, Blue Ridge Summit: Multilingual Matters, pp. 142–168. ISBN: 978-1-78892-187-9. DOI: 10.21832/9781788921879-009. URL: https://www.degruyter.com/document/doi/10.21832/9781788921879-009/html (visited on 05/05/2024).

Pennycook, Gordon and David G. Rand (Jan. 2021). "Research note: Examining false beliefs about voter fraud in the wake of the 2020 Presidential Election". en-US. In: *Harvard Kennedy School Misinformation Review*. DOI: 10.37016/mr-2020-51. (Visited on 01/16/2022).

Peterson-Salahuddin, Chelsea (Jan. 2022). "Posting Back: Exploring Platformed Black Feminist Communities on Twitter and Instagram". en. In: *Social Media + Society* 8.1, p. 20563051211069051. ISSN: 2056-3051. DOI: 10.1177/20563051211069051. URL: https://doi.org/10.1177/20563051211069051 (visited on 06/04/2024).

Philipp-Muller, Aviva, Spike W. S. Lee, and Richard E. Petty (July 2022). "Why are people anti-science, and what can we do about it?" In: *Proceedings of the National Academy of Sciences* 119.30, e2120755119. DOI: 10.1073/pnas.2120755119. URL: https://www.pnas.org/doi/10.1073/pnas.2120755119 (visited on 08/11/2022).

Phillips, Whitney (2018). "The oxygen of amplification". In: *Data & Society* 22, pp. 1–128.

Pickard, Victor (Apr. 2021). "Unseeing propaganda: How communication scholars learned to love commercial media". en-US. In: *Harvard Kennedy School Misinformation Review*. DOI: 10.37016/mr-2020-66. URL: https://misinforeview.hks.harvard.edu/article/unseeing-propaganda-how-communication-scholars-learned-to-love-commercial-media/ (visited on 05/16/2024).

Pierri, Francesco, Luca Luceri, and Emilio Ferrara (2022). "How does Twitter account moderation work? Dynamics of account creation and suspension during major geopolitical events". In: *arXiv preprint arXiv:2209.07614*.

Pirro, Andrea L P (2023). "Far right: The significance of an umbrella concept". en. In: *Nations and Nationalism* 29.1, pp. 101–112. ISSN: 1469-8129. DOI: 10.1111/nana.12860. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/nana.12860 (visited on 05/08/2024).

Pöyry, Essi et al. (2021). "A call for authenticity: Audience responses to social media influencer endorsements in strategic communication". In: *Social media influencers in strategic communication*. Routledge, pp. 103–118. URL: https://www.taylorfrancis.com/chapters/edit/10.4324/9781003181286-6/call-authenticity-audience-responses-social-media-influencer-endorsements-strategic-communication-essi-p\%C3\%B6yry-matilde-pelkonen-emma-naumanen-salla-maaria-laaksonen (visited on 05/07/2024).

Proferes, Nicholas (Jan. 2017). "Information Flow Solipsism in an Exploratory Study of Beliefs About Twitter". en. In: *Social Media + Society* 3.1, p. 2056305117698493. ISSN: 2056-3051. DOI: 10.1177/2056305117698493. URL: https://doi.org/10.1177/2056305117698493 (visited on 04/24/2020).

Pugh, J. (2009). *What Is Radical Politics Today?* London, UNITED KINGDOM: Palgrave Macmillan UK. ISBN: 978-0-230-25114-4. URL: http://ebookcentral.proquest.com/lib/washington/detail.action?docID=555564 (visited on 05/08/2024).

Radonovich, Lewis J. et al. (Sept. 2019). "N95 Respirators vs Medical Masks for Preventing Influenza Among Health Care Personnel: A Randomized Clinical Trial". eng. In: *JAMA* 322.9, pp. 824–833. ISSN: 1538-3598. DOI: 10.1001/jama.2019.11645.

Raina MacIntyre, C., Valentina Costantino, and Arjun Chanmugam (Oct. 2021). "The use of face masks during vaccine roll-out in New YorkCity and impact on epidemic control". en. In: *Vaccine* 39.42, pp. 6296–6301. ISSN: 0264-410X. DOI: 10.1016/j.vaccine.2021.08.102. URL: https://www.sciencedirect.com/science/article/pii/S0264410X21011646 (visited on 12/26/2021).

Rancourt, Denis G. (June 2020). *Masks Don't Work: A Review of Science Relevant to COVID-19 Social Policy*. en. URL: https://www.rcreader.com/commentary/masks-dont-work-covid-a-review-of-science-relevant-to-covide-19-social-policy (visited on 11/02/2022).

Randall, K. et al. (Oct. 2021). "How did we get here: what are droplets and aerosols and how far do they go? A historical perspective on the transmission of respiratory infectious diseases". In: *Interface Focus* 11.6, p. 20210049. DOI: 10.1098/rsfs.2021.0049. URL: https://royalsocietypublishing.org/doi/10.1098/rsfs.2021.0049 (visited on 05/25/2023).

Rathje, Steve, Jay J. Van Bavel, and Sander van der Linden (June 2021). "Out-group animosity drives engagement on social media". In: *Proceedings of the National Academy of Sciences* 118.26, e2024292118. DOI: 10.1073/pnas.2024292118. URL: https://www.pnas.org/doi/full/10.1073/pnas.2024292118 (visited on 04/16/2024).

Reagle, Joseph (Sept. 2022). "Disguising Reddit sources and the efficacy of ethical research". en. In: *Ethics and Information Technology* 24.3, p. 41. ISSN: 1572-8439. DOI: 10.1007/s10676-022-09663-w. URL: https://doi.org/10.1007/s10676-022-09663-w (visited on 04/21/2024).

Rengasamy, Samy, Benjamin Eimer, and Ronald E. Shaffer (Oct. 2010). "Simple Respiratory Protection–Evaluation of the Filtration Performance of Cloth Masks and Common Fabric Materials Against 20–1000 nm Size Particles". In: *The Annals of Occupational Hygiene* 54.7, pp. 789–798. ISSN: 0003-4878. DOI: 10.1093/annhyg/meq044. URL: https://doi.org/10.1093/annhyg/meq044 (visited on 11/21/2022).

Reuters (May 2021). "53% of Republicans view Trump as true U.S. president -Reuters/Ipsos". en. In: *Reuters*. URL: https://www.reuters.com/world/us/53-republicans-view-trump-true-us-president-reutersipsos-2021-05-24/ (visited on 01/16/2022).

Rheingold, Howard (Sept. 1993). *The Virtual Community: Homesteading On The Electronic Frontier The Edge*. en. Basic Books. ISBN: 978-0-201-60870-0.

Riedl, Magdalena et al. (2021). "The Rise of Political Influencers–Perspectives on a Trend Towards Meaningful Content". In: *Frontiers in Communication* 6. ISSN: 2297-900X. URL: https://www.frontiersin.org/articles/10.3389/fcomm.2021.752656 (visited on 11/07/2023).

Riedl, Martin J., Josephine Lukito, and Samuel C. Woolley (Apr. 2023). "Political Influencers on Social Media: An Introduction". en. In: *Social Media + Society* 9.2, p. 205630512311779. ISSN: 2056-3051, 2056-3051. DOI: 10.1177/20563051231177938. URL: http://journals.sagepub.com/doi/10.1177/20563051231177938 (visited on 11/08/2023).

Robb, Michael B. (2020). *Teens and the News: The Influencers, Celebrities, and Platforms They Say Matter Most*. en. Tech. rep. Common Sense. URL: https://www.commonsensemedia.org/research/teens-and-the-news-the-influencers-celebrities-and-platforms-they-say-matter-most-2020 (visited on 05/11/2023).

Robertson, Claire E. et al. (May 2023). "Negativity drives online news consumption". en. In: *Nature Human Behaviour* 7.5, pp. 812–822. ISSN: 2397-3374. DOI: 10.1038/s41562-023-01538-4. URL: https://www.nature.com/articles/s41562-023-01538-4 (visited on 04/16/2024).

Rohlinger, Deana A. et al. (Jan. 2023). "Does the Musk Twitter Takeover Matter? Political Influencers, Their Arguments, and the Quality of Information They Share". en. In: *Socius* 9, p. 23780231231152193. ISSN: 2378-0231. DOI: 10.1177/23780231231152193. URL: https://doi.org/10.1177/23780231231152193 (visited on 05/07/2024).

Rojek, Chris (2001). *Celebrity*. Reaktion books. URL: https://books.google.com/books?hl=en\&lr=\&id=jJ0QdF0WSIwC\&oi=fnd\&pg=PA9\&dq=chris+rojek+celebrity\&ots=c877kXYkGg\&sig=Dnugh1WUmSkSYWEn3j1mSfOMGrQ (visited on 04/18/2024).

Romm, Tony and Elizabeth Dwoskin (Jan. 2021). "Twitter purged more than 70,000 accounts affiliated with QAnon following Capitol riot". en. In: *Washington Post*. URL: https://www.washingtonpost.com/technology/2021/01/11/trump-twitter-ban/ (visited on 08/03/2021).

Rosvall, Martin, Daniel Axelsson, and Carl T Bergstrom (2009). "The map equation". In: *The European Physical Journal Special Topics* 178.1, pp. 13–23.

Rosvall, Martin and Carl T. Bergstrom (Jan. 2010). "Mapping Change in Large Networks". en. In: *PLOS ONE* 5.1, e8694. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0008694. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0008694 (visited on 10/25/2022).

Rothut, Sophia et al. (May 2023). "Ambassadors of ideology: A conceptualization and computational investigation of far-right influencers, their networking structures, and communication practices". en. In: *New Media & Society*, p. 14614448231164409. ISSN: 1461-4448. DOI: 10.1177/14614448231164409. URL: https://doi.org/10.1177/14614448231164409 (visited on 04/20/2024).

Rydgren, Jens (Aug. 2007). "The Sociology of the Radical Right". en. In: *Annual Review of Sociology* 33.Volume 33, 2007, pp. 241–262. ISSN: 0360-0572, 1545-2115. DOI: 10.1146/annurev.soc.33.040406.131752. URL: https://www.annualreviews.org/content/journals/10.1146/annurev.soc.33.040406.131752 (visited on 05/08/2024).

Salehi, Niloufar et al. (Apr. 2023). "Sustained Harm Over Time and Space Limits the External Function of Online Counterpublics for American Muslims". In: *Proceedings of the ACM on Human-Computer Interaction* 7.CSCW1, 93:1–93:24. DOI: 10.1145/3579526. URL: https://dl.acm.org/doi/10.1145/3579526 (visited on 10/23/2023).

Sandoval, Marisol and Christian Fuchs (May 2010). "Towards a critical theory of alternative media". en. In: *Telematics and Informatics*. Community media - the long march 27.2, pp. 141–150. ISSN: 0736-5853. DOI: 10.1016/j.tele.2009.06.011. URL: https://www.sciencedirect.com/science/article/pii/S0736585309000410 (visited on 07/29/2022).

Scheufele, Dietram A. and Nicole M. Krause (Apr. 2019). "Science audiences, misinformation, and fake news". In: *Proceedings of the National Academy of Sciences* 116.16, pp. 7662–7669. DOI: 10.1073/pnas.1805871115. URL: https://www.pnas.org/doi/full/10.1073/pnas.1805871115 (visited on 12/18/2022).

Schulz, Hans-Jörg and Christophe Hurter (2013). "Grooming the hairball-how to tidy up network visualizations?" In: *INFOVIS 2013, IEEE Information Visualization Conference*.

Schumacher, Elizabeth (Feb. 2022). *Disclose.TV: English disinformation made in Germany*. en. URL: https://www.dw.com/en/disclosetv-english-disinformation-made-in-germany/a-60694332 (visited on 04/08/2024).

Schwedel, Heather (Dec. 2017). ""Dunking" Is Delicious Sport". en-US. In: *Slate*. ISSN: 1091-2339. URL: https://slate.com/technology/2017/12/dunking-is-delicious-and-also-probably-making-twitter-terrible.html (visited on 03/18/2024).

Scire, Sarah (Mar. 2024). *A window into Facebook closes as Meta sets a date to shut down CrowdTangle*. URL: https://www.niemanlab.org/2024/03/a-window-into-facebook-closes-as-meta-sets-a-date-to-shut-down-crowdtangle/ (visited on 04/16/2024).

Senft, Theresa M. (2008). *Camgirls: Celebrity and Community in the Age of Social Networks*. en. Peter Lang. ISBN: 978-0-8204-5694-2.

– (Feb. 2013). "Microcelebrity and the branded self". en. In: *A Companion to New Media Dynamics*. Ed. by John Hartley, Jean Burgess, and Axel Bruns. 1st ed. Wiley. ISBN: 978-1-4443-3224-7 978-1-118-32160-7. DOI: 10.1002/9781118321607. URL: https://onlinelibrary.wiley.com/doi/book/10.1002/9781118321607 (visited on 11/07/2023).

Shaw Jr, WM (1985). "Critical thresholds in co-citation graphs". In: *Journal of the American Society for Information Science* 36.1, pp. 38–43.

Shearer, Elisa (Jan. 2021). *More than eight-in-ten Americans get news from digital devices*. en-US. URL: https://www.pewresearch.org/short-reads/2021/01/12/more-than-eight-in-ten-americans-get-news-from-digital-devices/ (visited on 11/09/2023).

Shugars, Sarah et al. (Apr. 2021). "Pandemics, Protests, and Publics: Demographic Activity and Engagement on Twitter in 2020". en. In: *Journal of Quantitative Description: Digital Media* 1. ISSN: 2673-8813. DOI: 10.51685/jqd.2021.002. URL: https://journalqd.org/article/view/2570 (visited on 05/05/2024).

Shuster, Stef M. (2021). *Trans medicine: The emergence and practice of treating gender*. NYU Press.

Simpson, Ellen, Andrew Hamann, and Bryan Semaan (Jan. 2022). "How to Tame "Your" Algorithm: LGBTQ+ Users' Domestication of TikTok". en. In: *Proceedings of the ACM on Human-Computer Interaction* 6.GROUP, pp. 1–27. ISSN: 2573-0142. DOI: 10.1145/3492841. URL: https://dl.acm.org/doi/10.1145/3492841 (visited on 06/04/2024).

Simpson, Ellen and Bryan Semaan (Apr. 2023). "Rethinking Creative Labor: A Sociotechnical Examination of Creativity & Creative Work on TikTok". en. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Hamburg Germany: ACM, pp. 1–16. ISBN: 978-1-4503-9421-5. DOI: 10.1145/3544548.3580649. URL: https://dl.acm.org/doi/10.1145/3544548.3580649 (visited on 04/20/2024).

Singleton, Marilyn M. (Sept. 2020). *Mask Facts*. en-US. URL: https://aapsonline.org/mask-facts/ (visited on 11/02/2022).

Skoric, Marko M., Jing Liu, and Kokil Jaidka (Apr. 2020). "Electoral and Public Opinion Forecasts with Social Media Data: A Meta-Analysis". en. In: *Information* 11.4, p. 187. ISSN: 2078-2489. DOI: 10.3390/info11040187. URL: https://www.mdpi.com/2078-2489/11/4/187 (visited on 05/08/2023).

Skurnik, Ian et al. (Mar. 2005). "How Warnings about False Claims Become Recommendations". In: *Journal of Consumer Research* 31.4, pp. 713–724. ISSN: 0093-5301. DOI: 10.1086/426605. URL: https://doi.org/10.1086/426605 (visited on 10/25/2022).

Small, Henry (1973). "Co-citation in the scientific literature: A new measure of the relationship between two documents". en. In: *Journal of the American Society for Information Science* 24.4, pp. 265–269. ISSN: 1097-4571. DOI: 10.1002/asi.4630240406. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.4630240406 (visited on 10/25/2022).

Smith, C. Estelle, Eduardo Nevarez, and Haiyi Zhu (Apr. 2020). "Disseminating Research News in HCI: Perceived Hazards, How-To's, and Opportunities for Innovation". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. New York, NY, USA: Association for Computing Machinery, pp. 1–13. ISBN: 978-1-4503-6708-0. DOI: 10.1145/3313831.3376744. URL: https://doi.org/10.1145/3313831.3376744 (visited on 06/10/2023).

Smith, Jeffrey D. et al. (May 2016). "Effectiveness of N95 respirators versus surgical masks in protecting health care workers from acute respiratory infection: a systematic review and meta-analysis". eng. In: *CMAJ: Canadian Medical Association journal = journal de l'Association medicale canadienne* 188.8, pp. 567–574. ISSN: 1488-2329. DOI: 10.1503/cmaj.150835.

Soden, Robert et al. (Oct. 2021). "Time for Historicism in CSCW: An Invitation". en. In: *Proceedings of the ACM on Human-Computer Interaction* 5.CSCW2, pp. 1–18. ISSN: 2573-0142. DOI: 10.1145/3479603. URL: https://dl.acm.org/doi/10.1145/3479603 (visited on 04/26/2024).

Soto-Vásquez, Arthur D. et al. (Dec. 2020). "COVID-19: Contextualizing Misinformation Flows in a US Latinx Border Community (Media and Communication During COVID-19)". In: *Howard Journal of Communications* 0.0, pp. 1–19. DOI: 10.1080/10646175.2020.1860839. (Visited on 09/27/2021).

Spade, Jane and Craig Willse (2000). "Confronting the limits of gay hate crimes activism: A radical critique". In: *Chicano-Latino L. Rev.* 21, p. 38. URL: https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/chiclat21\&section=7 (visited on 05/04/2024).

Squires, Catherine R. (2002). "Rethinking the black public sphere: An alternative vocabulary for multiple public spheres". In: *Communication theory* 12.4, pp. 446–468.

Starbird, Kate (May 2017). "Examining the Alternative Media Ecosystem Through the Production of Alternative Narratives of Mass Shooting Events on Twitter". en. In: *Proceedings of the International AAAI Conference on Web and Social Media* 11.1, pp. 230–239. (Visited on 08/30/2021).

Starbird, Kate, Ahmer Arif, and Tom Wilson (Nov. 2019). "Disinformation as Collaborative Work: Surfacing the Participatory Nature of Strategic Information Operations". In: *Proceedings of the ACM on Human-Computer Interaction* 3.CSCW, 127:1–127:26. DOI: 10.1145/3359229. URL: https://doi.org/10.1145/3359229 (visited on 05/06/2020).

Starbird, Kate et al. (June 2018). "Ecosystem or Echo-System? Exploring Content Sharing across Alternative Media Domains". en. In: *Proceedings of the International AAAI Conference on Web and Social Media* 12.1. ISSN: 2334-0770. URL: https://ojs.aaai.org/index.php/ICWSM/article/view/15009 (visited on 12/31/2021).

Steele, Catherine Knight (Oct. 2021). *Digital Black Feminism*. New York University Press. ISBN: 978-1-4798-0839-7. DOI: 10.18574/nyu/9781479808373.001.0001. URL: https://www.degruyter.com/document/doi/10.18574/nyu/9781479808373.001.0001/html (visited on 03/12/2024).

Stekelenburg, Aart van et al. (Dec. 2022). "Scientific-Consensus Communication About Contested Science: A Preregistered Meta-Analysis". en. In: *Psychological Science* 33.12, pp. 1989–2008. ISSN: 0956-7976. DOI: 10.1177/09567976221083219. URL: https://doi.org/10.1177/09567976221083219 (visited on 06/11/2023).

Stevenson, Shawn (2021). *Mask Facts*. en-US. URL: https://themodelhealthshow.com/maskfacts/ (visited on 11/02/2022).

Stewart, Leo G. and Emma S. Spiro (Apr. 2021). "Nobody Puts Redditor in a Binary: Digital Demography, Collective Identities, and Gender in a Subreddit Network". In: *Proceedings of the ACM on Human-Computer Interaction* 5.CSCW1, 8:1–8:31. DOI: 10.1145/3449082. (Visited on 08/30/2021).

Stokman, Frans N. and Patrick Doreian (1997). "Evolution of social networks: processes and principles". In: *Evolution of social networks*, pp. 233–250.

Stone, Brad (Sept. 2009). "Share the Moment and Spread the Wealth". en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2009/09/27/business/27ping.html (visited on 12/09/2022).

Strolovitch, Dara Z. (2023). *When Bad Things Happen to Privileged People: Race, Gender, and What Makes a Crisis in America*. University of Chicago Press. URL: https://books.google.com/books?hl=en\&lr=\&id=Q3u7EAAAQBAJ\&oi=fnd\&pg=PR7\&dq=bad+things+happen+to+privileged+people\&ots=_N5CJSVWU-\&sig=KGCfM_YgtmlVQ2rZmj-7vfP-mfw (visited on 05/08/2024).

Stromer-Galley, Jennifer et al. (Oct. 2021). "Political Messaging Over Time: A Comparison of US Presidential Candidate Facebook Posts and Tweets in 2016 and 2020". en. In: *Social Media + Society* 7.4, p. 20563051211063465. ISSN: 2056-3051. DOI: 10.1177/20563051211063465. URL: https://doi.org/10.1177/20563051211063465 (visited on 05/18/2023).

Swann, Ben (July 2020). *Why Face Masks DON'T Work, According To SCIENCE*. URL: https://www.youtube.com/watch?v=h8upEg-bEJ8 (visited on 12/01/2022).

Tajfel, H (Jan. 1982). "Social Psychology of Intergroup Relations". en. In: *Annual Review of Psychology* 33.1, pp. 1–39. ISSN: 0066-4308, 1545-2085. DOI: 10.1146/annurev.ps.33.020182.000245. URL: https://www.annualreviews.org/doi/10.1146/annurev.ps.33.020182.000245 (visited on 04/16/2024).

Tajfel, Henri et al. (1979). "An integrative theory of intergroup conflict". In: *Organizational identity: A reader* 56.65, pp. 9780203505984–16.

TallBear, Kim (2013). *Native American DNA: Tribal belonging and the false promise of genetic science*. U of Minnesota Press.

Tang, Jiliang et al. (Sept. 2017). "A Survey of Signed Network Mining in Social Media". en. In: *ACM Computing Surveys* 49.3, pp. 1–37. ISSN: 0360-0300, 1557-7341. DOI: 10.1145/2956185. URL: https://dl.acm.org/doi/10.1145/2956185 (visited on 04/14/2024).

TeBlunthuis, Nathan, Valerie Hase, and Chung-Hong Chan (2024). "Misclassification in Automated Content Analysis Causes Bias in Regression. Can We Fix It? Yes We Can!" In: *Communication Methods and Measures* 0.0, pp. 1–22. ISSN: 1931-2458. DOI: 10.1080/19312458.2023.2293713. URL: https://doi.org/10.1080/19312458.2023.2293713 (visited on 05/29/2024).

Thakur, Dhanaraj and DeVan Hankerson Madrigal (Oct. 2022). *An Unrepresentative Democracy: How Disinformation and Online Abuse Hinder Women of Color Political Candidates in the United States*. en-US. Tech. rep. Center for Democracy & Technology. URL: https://cdt.org/insights/an-unrepresentative-democracy-how-disinformation-and-online-abuse-hinder-women-of-color-political-candidates-in-the-united-states/ (visited on 09/28/2023).

The Grayzone - Bias and Credibility (Apr. 2024). en-US. URL: https://mediabiasfactcheck.com/the-grayzone/ (visited on 04/09/2024).

Thompson, Stuart A. (Dec. 2022). "Musk Lifted Bans for Thousands on Twitter. Here's What They're Tweeting." en-US. In: *The New York Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2022/12/22/technology/musk-twitter-bans.html (visited on 01/16/2023).

Thorbjørnsrud, Kjersti and Tine Ustad Figenschou (May 2022). "The Alarmed Citizen: Fear, Mistrust, and Alternative Media". In: *Journalism Practice* 16.5, pp. 1018–1035. ISSN: 1751-2786. DOI: 10.1080/17512786.2020.1825113. URL: https://doi.org/10.1080/17512786.2020.1825113 (visited on 05/02/2024).

Timberg, Craig, Elizabeth Dwoskin, and Reed Albergotti (2021). "Inside Facebook, Jan. 6 violence fueled anger, regret over missed warning signs". In: *The Washington Post*.

Tollefson, Jeff (Feb. 2021). "Tracking QAnon: how Trump turned conspiracy-theory research upside down". en. In: *Nature* 590.7845, pp. 192–193. DOI: 10.1038/d41586-021-00257-y. URL: https://www.nature.com/articles/d41586-021-00257-y (visited on 05/03/2021).

Torres-Lugo, Christopher, Kai-Cheng Yang, and Filippo Menczer (May 2022). "The Manufacture of Partisan Echo Chambers by Follow Train Abuse on Twitter". en. In: *Proceedings of the International AAAI Conference on Web and Social Media* 16, pp. 1017–1028. ISSN: 2334-0770. DOI: 10.1609/icwsm.v16i1.19354. URL: https://ojs.aaai.org/index.php/ICWSM/article/view/19354 (visited on 06/03/2024).

Traag, V. A. and Jeroen Bruggeman (Sept. 2009). "Community detection in networks with positive and negative links". eng. In: *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics* 80.3 Pt 2, p. 036115. ISSN: 1550-2376. DOI: 10.1103/PhysRevE.80.036115.

Traag, V. A., L. Waltman, and N. J. van Eck (Mar. 2019). "From Louvain to Leiden: guaranteeing well-connected communities". en. In: *Scientific Reports* 9.1, p. 5233. ISSN: 2045-2322. DOI: 10.1038/s41598-019-41695-z. URL: https://www.nature.com/articles/s41598-019-41695-z (visited on 12/16/2021).

Tremayne, Mark et al. (Oct. 2006). "Issue Publics on the Web: Applying Network Theory to the War Blogosphere". In: *Journal of Computer-Mediated Communication* 12.1, pp. 290–310. ISSN: 1083-6101. DOI: 10.1111/j.1083-6101.2006.00326.x. URL: https://doi.org/10.1111/j.1083-6101.2006.00326.x (visited on 05/16/2023).

Tromble, Rebekah (Jan. 2021). "Where Have All the Data Gone? A Critical Reflection on Academic Digital Research in the Post-API Age". en. In: *Social Media + Society* 7.1, p. 2056305121988929. ISSN: 2056-3051. DOI: 10.1177/2056305121988929. URL: https://doi.org/10.1177/2056305121988929 (visited on 04/22/2024).

Tromble, Rebekah, Andreas Storz, and Daniela Stockmann (Nov. 2017). *We Don't Know What We Don't Know: When and How the Use of Twitter's Public APIs Biases Scientific Inference*. en. SSRN Scholarly Paper. Rochester, NY. DOI: 10.2139/ssrn.3079927. URL: https://papers.ssrn.com/abstract=3079927 (visited on 04/22/2024).

Trujillo, Caleb M. and Tammy M. Long (Jan. 2018). "Document co-citation analysis to enhance transdisciplinary research". In: *Science Advances* 4.1, e1701130. DOI: 10.1126/sciadv.1701130. URL: https://www.science.org/doi/10.1126/sciadv.1701130 (visited on 02/16/2023).

Tucker Carlson Tonight (July 2021). *Tucker Carlson Tonight : FOXNEWSW : July 1, 2021 10:00pm-11:00pm PDT*. eng. URL: http://archive.org/details/FOXNEWSW_20210702_050000_Tucker_Carlson_Tonight (visited on 06/17/2022).

Twitter (Jan. 2021). *Permanent suspension of realDonaldTrump*. en. URL: https://blog.twitter.com/en_us/topics/company/2020/suspension (visited on 03/28/2022).

Tynes, Brendesha M. and Kimberly J. Mitchell (June 2014). "Black Youth Beyond the Digital Divide: Age and Gender Differences in Internet Use, Communication Patterns, and Victimization Experiences". en. In: *Journal of Black Psychology* 40.3, pp. 291–307. ISSN: 0095-7984. DOI: 10.1177/0095798413487555. URL: https://doi.org/10.1177/0095798413487555 (visited on 05/04/2024).

Tyson, Neil deGrasse (Mar. 2021). *Neil deGrasse Tyson on the Pandemic Year: Science Needs Better Marketing*. en-US. URL: https://www.wsj.com/articles/neil-degrasse-tyson-on-the-pandemic-year-science-needs-better-marketing-11616106660 (visited on 10/25/2022).

Unsafe (Mar. 2024). *Unsafe: Meta Fails to Moderate Extreme Anti-trans Hate Across Facebook, Instagram, and Threads*. en-US. Tech. rep. GLAAD. URL: https://glaad.org/smsi/report-meta-fails-to-moderate-extreme-anti-trans-hate-across-facebook-instagram-and-threads/ (visited on 05/01/2024).

Usher, Nikki, Jesse Holcomb, and Justin Littman (July 2018). "Twitter Makes It Worse: Political Journalists, Gendered Echo Chambers, and the Amplification of Gender Bias". en. In: *The International Journal of Press/Politics* 23.3, pp. 324–344. ISSN: 1940-1612. DOI: 10.1177/1940161218781254. URL: https://doi.org/10.1177/1940161218781254 (visited on 10/03/2023).

Valente, Thomas W. (July 2012). "Network Interventions". In: *Science* 337.6090, pp. 49–53. DOI: 10.1126/science.1217330. URL: https://www-science-org.offcampus.lib.washington.edu/doi/10.1126/science.1217330 (visited on 12/15/2022).

Varnelis, Kazys (2008). *Networked publics*. Mit Press.

Venturini, Tommaso, Mathieu Jacomy, and Pablo Jensen (Jan. 2021). "What do we see when we look at networks: Visual network analysis, relational ambiguity, and force-directed layouts". en. In: *Big Data & Society* 8.1, p. 20539517211018488. ISSN: 2053-9517. DOI: 10.1177/20539517211018488. URL: https://doi.org/10.1177/20539517211018488 (visited on 04/24/2024).

VersoBooks (Nov. 2023). *PollyWilkins jacobin*. qme. Tweet. URL: https://twitter.com/VersoBooks/status/1730271845935362104 (visited on 05/02/2024).

Villa-Cox, Ramon et al. (June 2020). *Stance in Replies and Quotes (SRQ): A New Dataset For Learning Stance in Twitter Conversations*. DOI: 10.48550/arXiv.2006.00691. URL: http://arxiv.org/abs/2006.00691 (visited on 04/03/2023).

Vogels, Emily A. et al. (2021). "Americans and 'cancel culture': Where some see calls for accountability, others see censorship, punishment". In: URL: https://policycommons.net/artifacts/1529284/americans-and-cancel-culture/2219059/ (visited on 04/17/2024).

Vreese, Claes de and Rebekah Tromble (May 2023). "The Data Abyss: How Lack of Data Access Leaves Research and Society in the Dark". In: *Political Communication* 40.3, pp. 356–360. ISSN: 1058-4609. DOI: 10.1080/10584609.2023.2207488. URL: https://doi.org/10.1080/10584609.2023.2207488 (visited on 04/22/2024).

Vrontis, Demetris et al. (2021). "Social media influencer marketing: A systematic review, integrative framework and future research agenda". en. In: *International Journal of Consumer Studies* 45.4, pp. 617–644. ISSN: 1470-6431. DOI: 10.1111/ijcs.12647. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/ijcs.12647 (visited on 05/07/2024).

Vynck, Gerrit De (Apr. 2024). "The AI hype bubble is deflating. Now comes the hard part." en-US. In: *Washington Post*. ISSN: 0190-8286. URL: https://www.washingtonpost.com/technology/2024/04/18/ai-bubble-hype-dying-money/ (visited on 05/01/2024).

Wang, Yu et al. (May 2020). "Reduction of secondary transmission of SARS-CoV-2 in households by face mask use, disinfection and social distancing: a cohort study in Beijing, China". en. In: *BMJ Global Health* 5.5, e002794. ISSN: 2059-7908. DOI: 10.1136/bmjgh-2020-002794. URL: https://gh.bmj.com/content/5/5/e002794 (visited on 12/27/2021).

Wanless, A and J Pamment (2019). "How Do You Define a Problem Like Influence?" In: *Journal of Information Warfare* 18.3, pp. 1–14. (Visited on 01/16/2022).

Ward, Isabella (Dec. 2023). "Disinformation Researchers Adjust to Restrictions From Platforms". en. In: *Bloomberg*. URL: https://www.bloomberg.com/news/articles/2023-12-07/social-media-researchers-struggle-to-track-disinformation-on-x-meta (visited on 04/16/2024).

Warzel, Charlie (Dec. 2022). "Elon Musk Is a Far-Right Activist". en. In: *The Atlantic*. URL: https://www.theatlantic.com/technology/archive/2022/12/elon-musk-twitter-far-right-activist/672436/ (visited on 11/05/2023).

Webster, James G. (Feb. 2011). "The Duality of Media: A Structurational Theory of Public Attention". In: *Communication Theory* 21.1, pp. 43–66. ISSN: 1050-3293. DOI: 10.1111/j.1468-2885.2010.01375.x. URL: https://doi.org/10.1111/j.1468-2885.2010.01375.x (visited on 05/29/2024).

Weigel, Moira (Jan. 2022). "The Authoritarian Personality 2.0". en. In: *Polity* 54.1, pp. 146–180. ISSN: 0032-3497, 1744-1684. DOI: 10.1086/717511. URL: https://www.journals.uchicago.edu/doi/10.1086/717511 (visited on 11/06/2023).

– (2023). "Hating Theory:"Cultural Marxism,""CRT," and the Power of Media Affects". In: URL: https://repositorio.consejodecomunicacion.gob.ec/handle/CONSEJO_REP/6670 (visited on 06/03/2024).

Weigel, Moira and Adina Gitomer (Apr. 2024). "Hate-sharing: A case study of its prevalence and impact on Gab". en. In: *New Media & Society*, p. 14614448241245349. ISSN: 1461-4448, 1461-7315. DOI: 10.1177/14614448241245349. URL: https://journals.sagepub.com/doi/10.1177/14614448241245349 (visited on 06/03/2024).

Weimann, Gabriel (1991). "The influentials: back to the concept of opinion leaders?" In: *Public Opinion Quarterly* 55.2, pp. 267–279.

Welbers, Kasper et al. (Nov. 2016). "News selection criteria in the digital age: Professional norms versus online audience metrics". en. In: *Journalism* 17.8, pp. 1037–1053. ISSN: 1464-8849. DOI: 10.1177/1464884915595474. URL: https://doi.org/10.1177/1464884915595474 (visited on 05/02/2024).

Whittaker, Meredith (Nov. 2021). "The steep cost of capture". en. In: *Interactions* 28.6, pp. 50–55. ISSN: 1072-5520, 1558-3449. DOI: 10.1145/3488666. URL: https://dl.acm.org/doi/10.1145/3488666 (visited on 04/12/2024).

Williams, Paige (Oct. 2020). "Inside the Lincoln Project's War Against Trump". en-US. In: *The New Yorker*. ISSN: 0028-792X. URL: https://www.newyorker.com/magazine/2020/10/12/inside-the-lincoln-projects-war-against-trump (visited on 04/09/2024).

Wojcieszak, Magdalena et al. (2022). "Most users do not follow political elites on Twitter; those who do show overwhelming preferences for ideological congruity". In: *Science advances* 8.39, eabn9418.

Wojcik, Stefan and Adam Hughes (2019). "Sizing up Twitter users". In: *PEW research center* 24, pp. 1–23. URL: https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2019/04/twitter_opinions_4_18_final_clean.pdf (visited on 05/05/2024).

Wolf, Thomas et al. (July 2020). *HuggingFace's Transformers: State-of-the-art Natural Language Processing*. URL: http://arxiv.org/abs/1910.03771 (visited on 04/05/2024).

Yang, Aimei et al. (May 2021). "The Influence of Interdependence in Networked Publics Spheres: How Community-Level Interactions Affect the Evolution of Topics in Online Discourse". In: *Journal of Computer-Mediated Communication* 26.3, pp. 148–166. ISSN: 1083-6101. DOI: 10.1093/jcmc/zmab002. URL: https://doi.org/10.1093/jcmc/zmab002 (visited on 05/16/2023).

Young, Meg, Michael Katell, and P.M. Krafft (June 2022). "Confronting Power and Corporate Capture at the FAccT Conference". In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '22. New York, NY, USA: Association for Computing Machinery, pp. 1375–1386. ISBN: 978-1-4503-9352-2. DOI: 10.1145/3531146.3533194. URL: https://dl.acm.org/doi/10.1145/3531146.3533194 (visited on 05/04/2024).

Zhang, Weiyu (Oct. 2016). "Social media and elections in Singapore: comparing 2011 and 2015". In: *Chinese Journal of Communication* 9.4, pp. 367–384. ISSN: 1754-4750. DOI: 10.1080/17544750.2016.1231129. URL: https://doi.org/10.1080/17544750.2016.1231129 (visited on 05/22/2023).

Zheng, Cheng et al. (2020). "Measuring time-sensitive and topic-specific influence in social networks with LSTM and self-attention". In: *IEEE Access* 8, pp. 82481–82492. URL: https://ieeexplore.ieee.org/abstract/document/9083959/ (visited on 04/13/2024).

Zolfaghar, Kiyana and Abdollah Aghaie (2010). "Mining trust and distrust relationships in social web applications". In: *Proceedings of the 2010 IEEE 6th international conference on intelligent computer communication and processing*. IEEE, pp. 73–80. URL: https://ieeexplore.ieee.org/abstract/document/5606460 (visited on 04/14/2024).