Social Media and Information Polarization: Amplifying Echoes or Extremes?

by

Julia Kamin

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Political Science)
in The University of Michigan
2019

Doctoral Committee:

Professor Nicholas Valentino, Chair
Assistant Professor Ceren Budak
Professor Scott Page
Professor Stuart Soroka

Julia Kamin

julkamin@umich.edu

ORCID iD: 0000-0003-1748-6935

Acknowledgments

Ten years ago during a period of professional soul searching I confessed to a friend I dreamed of writing a book about human behavior and democracy. When asked why I didn't just write it I protested I first needed a few years to read and think. She responded: "Have you ever heard of a PhD program?" To that friend I owe my first debt of gratitude for launching me in this great journey, the most intellectually fulfilling of my life. There are countless others to whom thanks are due. I broadened my understanding of American democracy thanks to Rob Mickey and Rick Hall and deepened my knowledge of social and political psychology under the guidance of Ted Brader and Phoebe Ellsworth. Cliff Lampe, Michael Heaney, Daniel Romero, Eytan Adar, and Rick Riolo all gave me the perspective and tools to study, explore and model complex online social networks. Innumerable unofficial mentors and supportive colleagues coaxed, encouraged and guided my thinking including Jennifer Chudy, Tim Ryan, Spencer Piston, Josh Pasek, Zander Furnas, Chris Skovron, Neil Lewis, Kentaro Toyama and, far from least, the two gentlemen who left no idea unchallenged and every aspiration encouraged, Fabian Neuner and Hakeem Jefferson. In the development of this dissertation, special thanks go to Skip Lupia for honing my research approach to studying Twitter and to my RA Matthew Stewart for setting me up to be a Twitter data vacuum. And then there is my perfect committee. Gratitude goes to Ceren Budak for forcing me to be precise in my assumptions and alerting me to inferential traps (a role I usually like to play). Stuart Soroka provided invaluable, careful feedback to multiple drafts making the work far more coherent and stronger than I could have hoped. Scott Page, whose mind and encouragement could inspire a thousand dissertations, is responsible for propelling my work in modelling diffusion on social networks which laid the foundation of this dissertation. And finally to my advisor. I hit the jackpot in getting Nicholas Valentino as my mentor and guide. Nick is the person to go to when your thoughts are hazy or you are struggling with a research design; his patience and sharp mind will guarantee you are left with clarity and direction every time. This work is as much a product of Nick's endless cajoling, critiquing, enthusiasm and

encouragement as it is mine. To him I owe the biggest debt, for making my ten-year old dream come true.

Table of Contents

Acknowledgments ii

List of Tables vi

List of Figures vii

List of Appendices xi

Abstract xii

Introduction 1

Part I: The Dynamics of Information Sorting and Extreming 27

Chapter 1. The Mechanics of Information Distribution 28

Chapter 2. Diffusion 38

Chapter 3. Homophily 74

Chapter 4. Curation 88

Chapter 5. Re-parametrizing the Model 110

Part II. A Theory of Information Sharing: Why Social Media Users are Motivated to Share Congruent and “Extreme” Information 117

Chapter 6: Building a Theory of Political Information Sharing 118

Chapter 7: Putting the Theory to the Test 140

Conclusion 184

Appendices 199

Bibliography 211

List of Tables

Table 1. Homophily in Offline Networks 84

Table 2. Examples of Online News Sites in Each Decile of the Bakshy Index 102

Table 3. Reasons for Sharing Political Information 204

List of Figures

Figure 1: The Author's Political Bubble on Facebook 12

Figure 2. Conceiving of Information Sorting and Information Extreming as Two Distinct Dimensions 20

Figure 3. A Slice of an Infinite Directed Network... 40

Figure 4. A Conservative Message Diffusing Toward a Sorting Equilibrium............................ 42

Figure 5. Sorting Equilibria across Networks of Different Levels of Homophily....................... 44

Figure 6. Diffusion's Effect on Information Sorting: The Starting Point Matters....................... 46

Figure 7. The Conditions in which Diffusion Leads to Information Sorting............................... 47

Figure 8. Conditions for Information Sorting across Networks with Varying Rates of Homophily 48

Figure 9. Example Random Binary Network.. 53

Figure 10. Example Binary Small World Network... 55

Figure 11. Example Continuous Random Graph 56

Figure 12. Diffusion in Agent-Based Models: Information Sorting at the Network Level......... 60

Figure 13. Diffusion in Agent-Based Models: A Selective Starting Point 61

Figure 14. Diffusion in Agent-Based Models: Information Sorting at the Node Level 62

Figure 15. Conceptualizing Binary Selective Curation 64

Figure 16. Conceptualizing Extreme Selective Curation.. 66

Figure 17. The Viral Edge of Extreme Messages 68

Figure 18. When Ideological News Will Spread 71

Figure 19. Two Ways to Calculate Network Homophily 81

Figure 20. Ideology of Twitter Users in Sample 94

Figure 21. De facto Selective Curation: The Association Between User Ideology and Ideology of their Tweets on Twitter 95

Figure 22. Ideology of Incoming and Outgoing Tweets 96

Figure 23. Ideological Difference Between What Users See and Retweet 98

Figure 24. Selective Curation Mediated by Selective Exposure 100

Figure 25. Aggregate Extreme Curation: Panelists in Sample 103

Figure 26. Aggregate Extreme Curation: Random Sample of Twitter Users 104

Figure 27. Signs of Extreme Curation on Facebook 105

Figure 28. More Signs of Extreme Curation on Facebook and Twitter 105

Figure 29. Average Extreme Curation 107

Figure 30. Modeling Information Sorting with Real Life Parameters 112

Figure 31. Modeling Extreme Curation with Real Life Parameters 113

Figure 32. Information Sorting on Facebook 115

Figure 33. An Example Set of Stories in the Survey 144

Figure 34. Selective Curation of Party Favorable News 148

Figure 35. Comparing Selective Curation and Selective Exposure 152

Figure 36. Selective Curation of Weakly vs Strongly Dogmatic News.................................... 159

Figure 37. A Post-Hoc Analysis: Dogma and Selective Curation 160

Figure 38. Example Stories: Superior Group Traits 162

Figure 39. Example Stories: In-party Credit 163

Figure 40. Example Stories: Outgroup Immorality 164

Figure 41. Extolling the Ingroup / Denigrating the Outgroup 165

Figure 42. Extolling the Ingroup: Superior Traits.. 166

Figure 43. Extolling the Ingroup: Ingroup Credit 167

Figure 44. Denigrating the Outgroup: Morality.. 168

Figure 45. Group Neutral Story for Republican Immorality / Racism 169

Figure 46. Group Neutral Story for Democrat Immorality / Racism.. 169

Figure 47. Morality - Comparing Democrats & Republicans.. 170

Figure 48. Effects of Priming Identity on Shifting the “Imagined Audience” I 175

Figure 49. Effects of Priming Identity on Shifting the “Imagined Audience” II....................... 176

Figure 50. Effects of Priming Identity on Selective Exposure 177

Figure 51. Effects of Priming Identity on Shifting “Imagined Audience” III 178

Figure 52. Affirming Away a Need for Inclusion.. 181

Figure 53. Filter Bubbles in the Public’s Mind.. 185

Figure 54. A Popular Visualization of Online Echo Chambers 186

Figure 55. Fake News Enters the Discussion 187

Figure 56. Information Sorting Equilibria: Reprise 189

Figure 57. Diffusion in Agent-Based Models: Information Sorting at the Network Level 200

Figure 58. Diffusion in Agent-Based Models: Information Sorting at the Node Level 201

Figure 59. Reasons for Posting Political Information, Ordered from Top to Bottom 205

Figure 60. Reasons for Posting Political Information, Ordered from Top to Bottom, Among those who Post Political Information Themselves 207

Figure 61. Association between User Traits and Ideological Curation 210

List of Appendices

Appendix I: Agent-Based Simulations Across Homophily Levels…………………………..….. 200

Appendix II: Self-report and Naïve Theories of Selective Curation………………………….. 203

Appendix III: Traits that Predict Selective Curation……………….……………………………. 208

Abstract

Social media sites are widely believed to solidify our echo chambers by sorting the political information liberal and conservative Americans are exposed to, leaving partisans and ideologues in information bubbles where their beliefs are confirmed and their views go uncontested. This dissertation challenges that common understanding of information polarization on social media and instead proposes that social media polarizes not by exposing users to congruent information but rather to extreme information. In examining the potential mechanisms of information sorting on social media – the flow of information in homophilous networks and users' choices in who they connect to and what to share – I find that social media sites are poorly equipped to filter out challenging information. Indeed, they may be better designed to expose users to more ideologically diverse information than those users encounter on other media on and offline. Rather than solidify our echo chambers, I propose that social media sites polarize our information environments by amplifying the prevalence of what I term "extreme" news - news that is ideologically dogmatic, emotionally-evocative and tribal in nature. Like "information sorting" and the creation of echo chambers, "information extreming" can have damaging consequences for civil society and democracy; the exposure to extreme information from both sides of the political spectrum leaves partisans and ideologues more certain in their correctness and in the illegitimacy of their political opponents. In the second part of the dissertation I examine one of the levers of information polarization – social media users' curation biases – and present a theoretical framework to explain users' motivations for sharing confirming and extreme information. That framework proposes that users are motivated to both signal solidarity with their political groups and, in times of threat, to rally against outgroups. Online experimental tests provide confirmation for elements of the framework, but much work is left to be done to understand what drives users to share congruent and extreme information.

Introduction

On the morning of November 9th, 2016, liberal America woke up both defeated and confused: defeated after having lost a hard-fought election battle; confused by the fact that half of their fellow Americans had voted for a man they believed to be indisputably unfit for office.

In the face of this bewildering result, liberals began to publicly wonder if they were dangerously out of touch. Perhaps they were living in bubbles that were blind to the lives of and views of large swaths of America? A Michael Moore Facebook post that chastised his followers to “stop saying they are ‘stunned’ and ‘shocked’” went viral: “What you mean to say is that you were in a bubble and weren’t paying attention to your fellow Americans and their despair” (Moore, 2016). Liberal elites asked what they should do to get out of their ideologically homogenous worlds. A *New York Times* op-ed “Life Outside the Liberal Bubble” (Vance, 2016), published the day after the election, provided a window into the mindsets of conservative working-class whites, and sales of the author’s book, *Hillbilly Elegy*, shot up soon after - as did Google searches for “liberal bubble” (Google Trends). Saturday Night Live didn’t miss the opportunity to have fun with their audience’s newly found liberal angst, advertising apartments in “The Bubble”: a new development where residents would find “a community of like-minded free thinkers....and no one else” (SNL, 2016).

Liberals also found a culprit for their cozy bubbles: social media. By giving Americans a place to connect and share information, social media sites like Facebook and Twitter were blamed for placing ideologues in information chambers where only their views are echoed back to them. So clear was social media’s role that if you were to search “echo chambers” in 2016, the top links Google would serve up included “The reason your feed became an echo chamber” (NPR Staff, 2016), “Confirmed: Echo chambers exist on social media” (Emba, 2016) and “Blame the echo chamber on Facebook: But blame yourself too" (Hosanakar, 2016). (The last, at least, acknowledged social media was not solely responsible.)[1]

Bubble making was not the only crime pinned on social media; social media sites were also faulted for distorting the 2016 election campaign by promoting "fake news." A cottage industry of sites had sprouted up in during the election, the owners of which weren't necessarily in it for political reasons (a few ran both liberal and conservative sites) but had become masters of fabricating stories ripe for the clicking. Trump supporters were fed headlines such as "Pope Francis Shocks World, Endorses Donald Trump for President" while Democrats could feast on "Rupaul Claims Trump Touched Him Inappropriately in the 1990s." And click they did. In the months before the election the top twenty fake news stories – which include the two above - had more "engagements" (likes, comments, shares) on Facebook than the top news stories from legitimate news sites (Silverman, 2016).[2]

While fake news made political headlines in 2016, the popularity of fabricated stories was part of a broader – though less publicly discussed - phenomenon on social media; the promotion of news that aims to push our partisan and ideological buttons. A few public thinkers took note. Jonathan Haidt warned of social media's capacity to expose users to a "constant stream of unbelievable outrages" (Haidt, 2016). Barack Obama similarly cautioned about social media's ability to spread information that paints "the opposition in wildly negative light without rebuttal" (Remnick, 2016). But by and large the spread of what I call "extreme" information – information that is ideologically dogmatic, emotional and tribal – has received less attention than echo chambers. Yet, as I will argue in this dissertation, the "extreming" of content in our social media feeds should cause us equal – if not greater – concern.

---

[1] Liberals are not alone in being in information bubbles; as will be discussed later in this introduction, left and right have increasingly been able to shield themselves in echo chambers. In 2016, however, few conservatives were publicly decrying those bubbles.

[2] Although 2016 certainly saw a surge in fake news, the extent to which fake news stories were viewed - let alone affected the election - is still an open question. One early study (Allcott & Gentzkow, 2017), for example, suggests that the average American only saw 1 or 2 fake news stories during the election cycle.

The two phenomena - information bubbles and information extreming – both polarize our information environments, but in different ways. Information bubbles sort political information between liberals and conservatives, leaving left and right in distinct information worlds, where their prior beliefs are left unchallenged and individuals are never forced to engage with those who disagree. Because the primary mechanism here is the ability to sort left from right information, I call this phenomenon "information sorting" – yet I will often use the terms "echo chamber" and "information bubble" to describe the result. "Information extreming," by contrast, exposes individuals to shocking and incendiary information. It polarizes by presenting information that represents the most strident and tribal views of each political side; rather than making us complacent with a general liberal or conservative point of view, exposure to extreme information makes us more dogged in our beliefs, including the conviction that our political opponents are unbalanced, unlikable, and untrustworthy.

This dissertation is an investigation of both phenomena – which I explore on two levels. First I look at social network sites as a whole, asking whether they do, in fact, exacerbate information sorting and information extreming. I then examine the psychology of social media users, to explain what may drive their choices to share political information that is congruent or extreme.

I offer an answer to the first question by unpacking the mechanics of social media to see if and how it has the capacity to polarize our information environments. While it may seem self-evident that social media is indeed designed to amplify echo chambers, I question that premise. Social networks are emergent systems and, as such, may not behave as our intuitions would expect. Social media distributes information through webs of individuals, each of whom make independent decisions about whom to connect to and what information to share or withhold. Those individual decisions, in turn, result in distribution patterns that are difficult to predict by using intuition or simple algorithmic models. To understand social media's capacity to polarize information, it is necessary to examine individuals' decisions about whom to connect to and what to share, but also to grasp how those micro decisions result in macro patterns of information distribution.

What I find challenges much of the common wisdom about social media. Social media sites do polarize our information environments, but not by sorting us into information bubbles. In fact, social media has the potential to expose us to a wider diversity of views and information than we would otherwise see in our daily lives. But the potential to break down echo chambers is not necessarily news to cheer, because while exposing us to information from our political opponents social media *is* well tuned to make sure the challenging information we see is extreme. The dynamics of what we share coupled with how we connect on social media leads to the amplification of extreme news – from left and right. Our bubbles are punctured with information from the other side on social media, but what seeps through is the most incendiary information – emotional, dogmatic and tribal. That is troubling for democracy.

While the first part of this dissertation looks at the global mechanisms of sorting and extreming, in Part II I zero in on one of those mechanisms – users' propensity to share news stories that reflect their views and are extreme in nature – to try to understand what drives those behaviors. I propose a theory of political information sharing founded on work in the disciplines of Social and Political Psychology, Anthropology, Sociology and Communications. That theory posits that social media users are motivated by two goals - to project an image of themselves to their political group that secures their inclusion and status in that group, and to rally their political groups during moments of perceived threat. In a series of experimental studies, I discover preliminary evidence that those goals will drive users to post news that is both consonant with their political views and that have the hallmarks of extreme news – dogmatism, emotional arousal, and tribalism.

Why Care about Information Sorting and Extreming

Democratic theorists and political psychologists give us several reasons to be concerned about the negative impacts on civic society of both information sorting and of information extreming.

By inhibiting the exchange of ideas and viewpoints across political groups, echo chambers stunt the growth of knowledge, stymie political compromise and even threaten the legitimacy of democracy itself. Societies thrive when individuals are exposed to a broad array of perspectives (Page, 2008). No individual or group has a monopoly on truth, but as J.S. Mill first argued societies are able to move toward that goal via through the “collision of adverse opinions” (Mill, 1869). A public marketplace of ideas, to use Holmes’ famous phrase (Abrams v. US, 1919), not only results in greater knowledge but is essential to a functioning democracy. When competing groups are not exposed to broad sets of information and form distinct perceptions of “truth,” it is difficult to find common ground and reach compromise (Sunstein, 2009). Potentially more destabilizing to civic society, when groups are unable to share their viewpoints the legitimacy of a democracy is threatened. Not all political groups can be electoral winners, but by exposing citizens to the reasons behind all parties’ preferences, public debate can legitimize those winners - and, by extension, legitimize our democratic institutions (Manin, 1987). Echo chambers, by contrast, not only cocoon us from (at least a portion of) the truth, but also prevent us from hearing our opponents’ reasoning and so leave us to imagine they are unreasonable or even dangerous.

Work in political psychology backs much of this normative view of information sorting. Even with the many cognitive and motivated biases known to prevent accurate processing of political information (Kunda, 1990; Lord et al, 1979; Taber & Lodge, 2006), citizens have at least some capacity to learn about which parties, candidates and policies best represent their preferences when they are exposed to information (Barabas & Jerit, 2009; Lupia, 1994; Redlawsk et al, 2010). Echo chambers, by narrowing our access to information, limit our ability to act as knowledgeable voters. Likewise in line with deliberative theorists’ claims, exposure to opposing views can promote tolerance for those with whom we disagree (Ben-Nun Bloom & Levitan, 2011; Erisen & Erisen, 2012; Mutz, 2002b; Mutz & Mondak, 2006) and even dampen hostility towards our political opponents (Garrett et al, 2014; Parsons, 2010).

Yet when those views are extreme as I define the term – dogmatic, emotion-laden and tribal - the outcomes of information exchange can create their own set of problems for civic society. Instead of promoting learning and awareness, exposure to the extreme versions of the views of our political opponents may further entrench us in our beliefs (Nelson et al, 2011). We may also infer that espousers of those extreme views must be morally corrupt or intellectually inferior (Ward et al, 1997), and make the leap to conclude that all of their co-ideologues are likewise not to be trusted (Iyengar et al, 2012; Lelkes et al, 2017; Levundusky, 2013). In short, when counter-ideologues share extreme views, instead of building greater tolerance and awareness between political groups, they may nurture distrust and animosity (Bail et al, 2018; Settle, 2016).

For a democratic society to be successful, then, it is not only important to ensure that ideas and beliefs are shared between liberals and conservatives, but that the views expressed are of a nature to inspire understanding rather than animus.

There are arguments, to be fair, that the effects of information sorting and extreme information are not all anti-democratic. Echo chambers may make us more certain and dogmatic in our beliefs, but such certainty also leads citizens to be engaged in the political process; ambivalence and moderation, in contrast, dampens political fervor, making citizens less likely to run to the polls, let alone to the next protest (Mutz, 2002b). If extreme information fuels animosity toward political opponents, it could similarly energize the electorate and increase political participation. Such animosity, moreover, may be warranted and even morally necessary; if it is directed at political foes who pose genuine threats to civil society deep animosity and distrust may be what are required to defeat those threats. Likewise, civil discourse and the tolerance and understanding it hopes to engender may seem like virtues of a democratic society, but civility can also be a guise to uphold the status quo and delegitimize the “uncivil” actions of oppressed groups (Mendelberg, 2002).

Those arguments have strong merits and I by no means take the position that all citizens should be ever openminded on all topics or that voters should not have strong political views or a zeal for defeating their political foes. Healthy democracies need debate, strong political opponents and engaged, fired up citizens. A political society in which all citizens are ambivalent and apathetic can be as dangerous as one in which all are fixed in their beliefs and ruled by inter-group animosity. As with all things, a balance is necessary. Strong views motivate action and debate, but when they become rigid and dogmatic a society stops learning. Partisan loyalty likewise organizes and motivates political participation, but when partisans view each other as fundamentally corrupt and beyond redemption, at best politics becomes zero-sum, with every political battle having only big winners and losers. At worst disagreements cannot be resolved through non-violent means and democracy breaks down.[3]

Information Sorting before Social Media

While the 2016 election season raised alarms about ideological echo chambers, the phenomenon is not new. Information sorting goes back at least to our founding, when many Americans had no choice but to live in political bubbles. Our early presses were partisan by design. While some news consumers might have been able to get a balanced information diet by, for example, purchasing both the Federalists' and Antifederalists' papers, presumably many opted for one partisan rag or the other, if they were lucky enough to live in an area where they even had a choice (Schudson, 1981). The concept of balanced, objective journalism didn't really take hold until the mid-20th century, and in that Golden Era of news there was enough diversity in the information environment - from political campaigns to chatty neighbors - to allow liberals and conservatives to pick and choose their information diets. In 1944, sociologists Lazarsfeld and colleagues quantified that divergence for the first time, observing that only 25% of Democrats and Republicans were exposed to proportionately more cross-partisan information than news from their own side (Lazarsfeld et al, 1944). Since that seminal work, dozens of studies have confirmed that Americans on the left and right live in information worlds that more often than not reflect their political views (Gentzkow & Shapiro, 2011; Sears & Freedman, 1967).

The causes of that information divide, however, have not been as easily observed and remain a matter of dispute. Lazarsfeld and colleague's groundbreaking study proposed two explanations for why political groups would end up in distinct information environments with little overlap.

---

[3] The idea that partisanship in the US could devolve into physical violence may have seemed far-fetched a couple of decades ago, but as signs of affective partisan attachment grow, as does the association between political and racial identities (Iyengar et al, 2012; Mason & Wronski, 2018), the future of Democrats and Republicans as Tutsis and Hutus or Serbs and Bosnians may not just be dystopian fantasy.

For one, voters have “psychological predispositions” and “somehow contrive to select out the passing stream of stimuli those by which they are more inclined to be persuaded” (p. 82). That is, they “selectively expose” themselves to information that aligns with their views. But what may play a larger role than individual inclinations is a voter’s likely existence in an environment that “sifts the propaganda” *for him* (p. 81). A voter need not look for politically aligned information to be in a bubble; his environment may create the bubble for him. Conservative rural dwellers, for example are more likely to read “farm journals,” which happen to lean Republican. City workers, conversely, are more apt to hear talk “from fellow-workers who are pro-labor and pro-Democratic” (p. 81). The authors also proposed that this congruent information reaches voters through a “two-step flow” from “opinion leaders” who first absorb information from the news media and political actors and then relay that information in their like-minded communities (p. 151).

Lazarsfeld and colleagues presented two puzzles regarding media consumption and social networks that remain relevant to our understanding of echo chambers today – and which social scientists still struggle to fully disentangle. One is understanding why people tend to surround themselves with like-minded others, a phenomenon social scientists call “homophily” (McPherson et al, 2001). Scholars know that the reason our friends, neighbors and co-workers do such a good job in sifting political information for us is because they, by and large, share our political views. What is unknown is exactly *why* they share our views. We suspect it is not primarily because we choose our friends because of their politics (Lazer et al, 2010).[4] Three other forces, instead, are likely more at play. First, we tend to encounter people who share our values and beliefs because we are drawn to the same places and institutions – our church, our job, our organic café – that reflect those values and beliefs. To the extent that those preferences are associated with political leanings, we will naturally find ourselves surrounded with more co-ideologues than not. We also tend to like people who are like us - who share our interests, habits, and way of thinking. Likewise, because those preferences often correlate with political

[4] Although we may choose our mates that way (Huber & Malhotra, 2013). Work in “affective polarization” also suggests Americans may increasingly eschew friendships with counter-partisans (Iyengar et al, 2012).

preferences, our friends will more often than not be politically aligned with us. Finally, homophily is reinforced because we take on the preferences and views of those around us; disagreement is uncomfortable and conformity has its benefits (Huckfeldt & Sprague, 1987; Lazer et al, 2010; McPherson et al, 2001; Sears & Freedman, 1967).

The other puzzle, which has been more the focus of political scientists, is the degree to which our political information is sorted via our environment or, instead, by our active choices. As pointed out by scholars Sears and Freedman (1967), the mere observation that voters consume information that aligns with their beliefs (what they call "de facto" selective exposure) does not mean they are actively seeking that information ("active" selective exposure).

At the time of those early works, there was relatively little room for active selective exposure to flourish compared to today. Before the 1980s, news consumers would be hard pressed to actively insulate themselves from opposing views. The federal "Fairness Doctrine" required that broadcasters on radio and TV present both sides of important issues. Although print media had no such restrictions, for the most part Americans had (and still have) one local newspaper to choose from, and that publication likely took a fairly middle of the road approach so as to not alienate portions of its readership (Prior, 2013). In such a media environment, although it would be possible to actively build an ideological news bubble, it would require considerable effort; not only would a consumer have to tune out all broadcast media, she would need to subscribe to a national weekly like The Nation or The National Review - or be lucky enough to live in a large enough city that could support multiple newspapers with distinct political viewpoints (Bennet & Iyengar, 2008; Katz & Lazarsfeld, 1955, Schudson, 1981; Sears & Freedman, 1967).[5] To the extent that information sorting existed in the mid-20th century it was likely due to incidental selective exposure, receiving information that trickled through Lazarsfeld's "opinion leaders."

Since the late 1980s, however, media choices have multiplied and diversified – creating an environment where active selective exposure can blossom. The Fairness Doctrine was revoked in

[5] That is not to say all Americans consumed the same information; depending on where they lived their local paper might have a left or right slant.

1987, giving broadcasters permission to present only one side of the story. Radio, in particular, ran with this option, leading to the rise of Rush Limbaugh and a long menu of other ideological (mostly conservative) talk shows. Cable news, although never restricted by broadcast regulations, built up large enough audiences to begin diversifying in the 1990's, birthing Fox and MSNBC (though MSNBC took a decade to become a safe space for liberals). The ideological fragmentation of radio and cable news would later be dwarfed by the explosion of bloggers and news aggregators and sites on the internet. By the early 2000's, online readers could get all their news mediated by Matt Drudge, Markos Moulitsas or countless other partisan online news sources and aggregators. Whereas in 1980 it would have taken great effort to build an ideological news bubble, 25 years later a conservative could tune in exclusively to Alex Jones, Fox and The Drudge Report while a liberal could likewise insulate himself in the cozy world of NPR, MSNBC and HuffPo.

Social commentators took note of the capacity for echo chambers to flourish online. Negroponte (1996) called the cozy internet environments we create the "Daily Me," and Sunstein (2009) warned people would use these enclaves to insulate themselves from counter-attitudinal messages. And they might not do so purely by choice; as Pariser (2011) suggested in the book that gave rise to the term "filter bubble," online algorithms would further segment our online news environment, pushing us into more and more tailored information bubbles.

Despite these warnings, research only weakly confirms that a fragmented media – on cable, talk radio and the internet – allows us to silo ourselves in congenial news environments. It is not even clear if consumers actively seek - or "selectively expose" themselves to – politically congruent information. Studies do consistently find that in experimental settings, when subjects have a choice on what news to consume, ideologues and partisans tend to prefer information that aligns with their views, although they don't show an aversion to uncongenial information and will even seek out information counter to their views under certain circumstances (for reviews see Frey, 1986; Hart et al, 2009). Whether news consumers actively select congruent information outside of the social scientist's lab, however, is harder to examine; again as Sears and Freedman (1967) point out, when you observe a liberal reading the New York Times and listening to NPR, it's difficult to say whether a) he has sought out those liberal sources, b) he's consuming them because that’s what he grew up with and what all his friends consume – or c) he is liberal because all he knows are the points of view of New York Times and NPR.

Stroud (2008), in a rare panel study that attempts to tease out what leads us to be exposed to predominantly congruent information, provides some evidence that “active selective exposure” is partly responsible. By collecting consumers’ ideological leanings and news preferences over time – and during an election season when presumably their interest in political news would rise – Stroud’s findings suggest that consumers’ ideological leanings do, in part, drive their consumption. Also in line with Daily Me predictions, she found such evidence only when looking at consumption of radio, cable and online news; ideology did not have a noticeable effect on consumers’ newspaper selection.

Stroud’s cross-media finding is partly seconded by perhaps the most comprehensive cross-media study of selective exposure. Gentzkow and Shapiro (2011) look at differences between consumption patterns of liberals and conservatives across print, television and the internet and also compare those divides to how segregated we are in face-to-face interactions. While they find that ideologues “isolate” themselves more on cable and the internet than they do when watching broadcast news, they also find that consumers are more divided in face-to-face interactions and in the national newspapers they read. Perhaps more surprising, consumers are farther apart in their choices of which magazines and local news to consume than on what to watch on cable.[6]

Information Sorting on Social Media

Social media is often presumed to continue the bubble-making trend set by radio, cable and the internet. Although thought by some early internet optimists to be an antidote to political information bubbles (Farrell, 2012; Himelboim et al, 2013), conventional wisdom today is that

---

[6] While Stroud and Gentzkow provide evidence that a fractured media leads to selective news exposure, perhaps the greater effect of media diversity, as shown by Prior (2005), is that it leads many consumers to select out of the news audience entirely.

social media further filters out counter-attitudinal voices, trapping us in echo chambers of like-minded folks – to the detriment of democratic deliberation and civic society. As mentioned above, Google searches for “echo chambers” in 2016 would primarily surface articles about social media. If you were to consult the Encyclopedia of Social Media and Politics it would likewise confirm that in using “social network sites, personal Web logs, and other Web platforms that encourage audience interaction, individuals strengthen the echo chambers by reposting media content” (Bor, 2014). To remind us of our walled informational worlds – and let us peek over to the other side - The Wall Street Journal keeps a running record of the differences between a “red” and “blue” feed on Facebook (WSJ). Likewise, Facebook apps have been developed to let users scan their “friendverses” and see how deep in the bubble they are. Here, for example, is this author’s, admittedly very blue, bubble:

Figure 1*: The Author’s Political Bubble on Facebook*

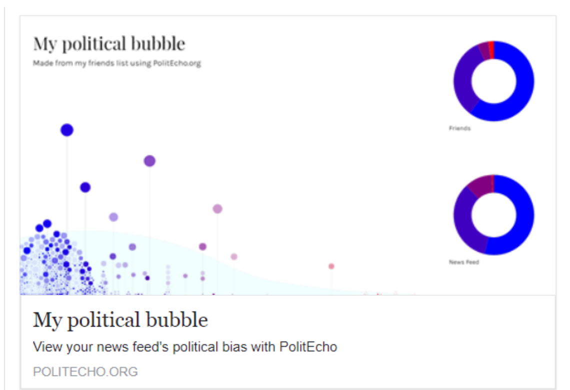

Image created by Politecho.org in 2016, using author’s Facebook data. Politecho’s app lets Facebook users chart the political diversity their friends – and how much of a Facebook bubble they presumably exist in.

The assumption that social media is designed to secure us in echo chambers may be popular because it makes intuitive sense. By being “social,” social media introduces three dynamics that seem to lend themselves to bubble formation. First and foremost, social media connects us to friends and, as in real life, those friends tend to share our views including our views on politics. Second, those friends determine what news we see; for the most part while news editors still create the news, our friends on social media are the ones who curate that news. It is natural to suspect that those friends would be biased curators, selectively sharing information that confirms their - and our - worldview. Finally, as in offline social networks, our friends tend to be friends with each other, embedding us inside clusters of like-minded ideologues; thus we may intuit that as friends re-share information they will act as filter upon filter further refining - and narrowing - the scope of information we are exposed to.[7]

In spite of these popularly held intuitions (see e.g. Sunstein, 2018; Manjoo, 2011), recent empirical research suggests that social media might not be the bubble maker often feared. True, studies that examine levels of cross-ideological exposure show that ideologues are mostly exposed to congruent information (Bakshy et al, 2015; Jungherr, 2014) but they do not tell us if these levels are more or less than we would expect to see in other media environments. One exception is a study by Flaxman and colleagues (2016) which shows that while social media does a poor job compared to search engines in leading us to diverse information, it does diversify our information environment more than if we were to go directly to news sources online. Another study (Barbera, 2014), which looks at information Twitter users are exposed to via social acquaintances compared to the news they receive from public news sources, finds similarly that those social connections do a better job of exposing us to diverse news.

Theorists likewise give us reason to imagine that social media could mitigate rather than exacerbate polarization. For one, homophily is far from absolute, so we should expect a non-negligible proportion of our connections to have differing political views (Huckfeldt et al, 2004). And, whereas before social media we might only talk politics with a small subset of our friends (and those closer to us), social media makes it easy - or unavoidable - to encounter the political views of a wider swath of acquaintances, including those we may have little offline contact with. Those “weak ties,” as network theorists call them, can expose us to viewpoints we would not have otherwise encountered (Brundidge, 2010). Yet another reason social media increases our

[7] There are other reasons (independent of these network mechanisms) we would reasonably think social media sites sort information. As highlighted by Eli Pariser in The Filter Bubble (2011), almost all social media sites use algorithms to filter the posts shown on users’ walls, selecting for information that users prefer to see based on their past choices and behavior. Advertisers – including political campaigns –also target messages to users based on profiles that social media sites construct (likewise based on user behavior). While these forces are certainly part of the information sorting story on social media, they are not the focus of this dissertation for various reasons which I discuss at length in the next chapter.

chances of bumping into counter-ideological information is its lack of topical barriers; whereas we might expect echo chambers to form if we could restrict political discourse to a forums about politics (in which me might select into like-minded groups), on social media politics will pop up regardless of the subject matter (Wojcieszak & Mutz, 2009).

The empirical evidence and theoretical insights above force us to question the popular narrative of social media sites as bubble makers and ask if social media is the echo chamber bogeyman it is commonly thought to be.

There are two ways to approach greater clarity. One is to expand on the findings of Flaxman et al (2016) and Barbera (2014) by conducting additional large-scale observational studies that compare the diversity of information that Americans are exposed to across different media sources, including print, TV, radio and water-cooler conversations. As mentioned above, such an approach was taken by Gentzkow and Shapiro (2010). However, while their research compared de facto selective exposure across several media environments, it stopped short of considering social media.

I take a second approach, instead breaking down the dynamics of social media and reexamining our intuitions about social media's ability to create echo chambers. While it may seem self-evident that social media is a perfect storm of forces designed to sort liberals and conservatives into distinct information worlds, no one has conclusively demonstrated that those forces exist, let alone measured their strength. I examine each in turn. I look at whether social media users are, indeed, selective curators who actively choose to post stories that reflect their political views. While accepting the conclusion from previous research that homophily is pervasive in online social networks (Bakshy et al, 2015; Barbera et al, 2015; Boutyline & Willer, 2017; Conover et al, 2011; Grabowicz et al, 2012; Hanna et al, 2011; Himelboim et al, 2013), I ask whether individuals in online networks are more homophilous than they are offline. Finally, I examine whether social networks – which embed us in chains of friendships, each of which is likely homophilous - amplify ideological homogeneity as information diffuses from user to user.

I find that several commonly held assumptions about social media's capacity to sort us into information bubbles are unsupported both theoretically and empirically. For one, social media's capacity to diffuse information through homophilous networks would not, in and of itself, logically produce information environments more homogenous than the group of friends, family, and associates to which we belong. Modelling diffusion in homophilous networks – i.e. networks in which users on average are friends with more co-ideologues than counter-ideologues - I show that our friends are limited in the degree to which they can filter our news for us; as information moves through social networks, it will never become fully sorted, but will rather be capped at the average level of homophily in the network.[8] Moreover, when our networks are less than 100% ideologically biased in the information they re-share, diffusion actually *diversifi*es our news environments, more so than if our friends did not re-share information at all.

Examining re-sharing behavior on the social network Twitter, I also find that users are surprisingly unbiased in the news they retweet (though their ideological biases show up in the news they choose to follow). Finally, relying on research from other scholars, I examine estimates of homophily online and in the real world and find little evidence to suggest our online networks are more homophilous than our networks offline. All told, I find social media is poorly equipped to filter out diverse views and may, perhaps, even expand the breadth of our exposure beyond what we would encounter in offline networks.

Polarizing by Extremes

The reader may be skeptical of the conclusion that social media is more bubble buster than builder. Such a claim runs counter to what seems an obvious truth about online social networks – that they are a polarizing force in today's world. I do not dispute that social media polarizes, but

[8] Throughout this work I will often speak of the relative homophily of networks. It is important to distinguish that by "homophily" in this respect I do not mean "homogeneity." It is possible, for example, to have a heterogeneous network – with 50% liberals and 50% conservatives – with high levels of homophily if, for example, individuals in that network almost exclusively connect to like-minded friends. A highly homogenous network with, say, 90% liberals can likewise have low homophily if liberals in that network make an effort to connect to the few conservatives that exist (which in turn would give those conservatives a surfeit of liberal friends).

I propose that instead of pulling left and right apart by sorting people into information bubbles, social media polarizes by proliferating “extreme” information on both sides.

By extreme information, I mean information - news articles, memes, videos - that is dogmatic, emotional and tribal. Extreme information is usually, though not always, generated by far left and right sources such as OccupyDemocrats or Breitbart news. It is also characteristic of much of the fake news that swirled through social media during the 2016 campaign. But extreme news is not necessarily – or even usually – false. Extreme information also does not necessarily advocate for extreme policies. What characterizes extreme news rather are three features: its tendency to be dogmatic in the support of an ideological view, its appeal to emotions such as anger and, perhaps most importantly, its tribal-like focus on political opponents.

While the above definition of extreme information is unique to this dissertation, I do not construct it out of whole cloth. Communication scholars and political scientists alike have noted and studied similar concepts – although often using different terms. Taylor (2017), for one, characterizes “extreme media” as media that has an “ideological point of view” and that “seeks to entertain” while it also “spurs emotions” such as anxiety and fear. In a review of “polarizing” media, Prior (2013) notes that such media uses “fervently populist or ideological rhetoric” and is “ideologically unambiguous” and “blatantly partisan.” Levendusky (2013) similarly describes this type of partisan media as creating “a coherent conservative or liberal interpretation” of the news, focusing on “criticism of the opposition,” and selecting stories that make the other side look “foolish and inept” as well as “hypocritical and duplicitous.” All three note partisan news outlets have a penchant for being “outrage-based” (Sobieraj and Berry 2013). Sobieraj and Berry’s definition of outrage discourse is worth reviewing in full:

> *“‘outrage’ as a particular form of political discourse involving efforts to provoke visceral responses (e.g., anger, righteousness, fear, moral indignation) from the audience through the use of overgeneralizations, sensationalism, misleading or patently inaccurate information, ad hominem attacks, and partial truths about opponents—opponents who may be individuals, organizations, or entire communities of interest (e.g., progressives*

or conservatives) or circumstance (e.g., immigrants, welfare recipients). Outrage sidesteps the messy nuances of complex political issues in favor of melodrama, misrepresentative exaggeration, mockery, and improbable forecasts of impending doom. Outrage speech is not as much about discussion as it is about verbal competition, akin to political theater with a scorecard."

Three common themes in the definitions above characterize extreme information:

- Dogmatism - Extreme news avoids "messy nuances" and instead offers a clear, unambiguous message, with a "coherent liberal or conservative interpretation."

- Emotion - Extreme news triggers powerful emotions by deploying sensationalist arguments and fervent language focusing on threats to cherished groups, society at large, or universal norms.

- Tribalism - Extreme news focuses on "criticism of the opposition," using "ad hominem attacks" against opponents. It is not about imparting information as much as participating in "competition" against the outgroup.

One theme that notably does not emerge from the definitions above is adherence to extreme *policy positions*. Extreme media may indeed advocate for far left or right policies, but that is not its defining, or even a necessary, feature. Extreme news can – and often will - push a centrist viewpoint, yet in a dogmatic, emotional or tribal manner. Take the example of trade tariffs, a topic in the news cycle as of this writing, and one not strongly associated with liberal or conservative views.[9] On a week in which China and the US both declare new tariffs, the New York Times headlines a straight journalistic "Trump Says He Will Raise Existing Tariffs on Chinese Goods to 30%" and Fox offers an equally tame "US and China announce new tariffs." Breitbart, however, opts for "Donald Trump Strikes Back at China with Higher Tariffs," a headline which paints the leader of its political group in a strong, even heroic light. Huffington

[9] Indeed, left and right both have their share of free trade advocates and protectionists.

Post's headline, "Trade War Tense - Markets Roiled - Recession Fears," goes further in its extreme tone by emphasizing threat and fear which, implicitly understood by any Huffington Post reader, is created by the opposing tribe.[10] Breitbart and Huffington Post both manage to take a policy neutral story and turn it into an extreme event.

One way we can think of the distinction between the definition of extreme news above - one of unambiguous, visceral, devotion to the tribe - as distinct from thinking of extreme news that might preach far left and right policies, is as a parallel to the distinction between traditional views of polarization and "affective polarization." Political scientists who quantify and track the degree of polarization in the US have historically focused on how far left and right legislators and voters are divided on policy preferences (Hetherington, 2009). That work, notably articulated by Fiorina et al (2005), shows that Americans, even staunch partisans, tend to be closer on core policy preferences than frequently imagined. Iyengar et al (2012) and other political scientists, however, point out that while Americans may not polarize on policy issues, they are becoming more deeply divided in their feelings toward each other.

That work on affective polarization echoes the work of an earlier theorist, Eric Hoffer, on "true believers" who (to use Fiorina's words) "place more weight on symbols (dubbed 'principles'), reject what appear to be reasonable compromises, draw bright lines where many people see only fuzzy distinctions, and label those who disagree with them as enemies" (Fiorina, 1999). Like affective partisanship or the rise of true believers, when we speak of extreme news, it is the emotional, symbolic and tribal aspect of political news, rather than a division of policy preferences, that is the defining feature.

Information extreming – the process of amplifying extreme viewpoints - may seem difficult to differentiate from information sorting, but they are conceptually distinct. When information is sorted, ideologues are placed in bubbles surrounded by relatively homogenous information that affirms their political views. From the perspective of the network as a whole, sorting results in

[10] Headlines taken from nytimes.com, foxnews.com, breitbart.com and huffingtonpost.com on August 23, 2019.

liberal information being concentrated among liberals and conservative information among conservatives. Information extreming, in contrast, results in more extreme - as compared to moderate - information *throughout* the network. At the individual level information extreming may leave liberals and conservatives alike exposed to diverse views, but a greater proportion of those views will be extreme (dogmatic, divisive, emotional) as opposed to moderate (nuanced, non-tribal, reasoned).

To see how information sorting and extreming lie on distinct dimensions we can imagine two information universes. In one, there are only two sources of information, NPR and the Wall Street Journal; liberals only listen to NPR while conservatives get their news exclusively from WSJ. This world, illustrated in Figure 2a, would be one of complete information sorting, but with relatively little extreme information. In another imaginary universe, visualized in Figure 2b, all liberals and conservatives might be equally exposed to left and right information, but the only two sources of information are Keith Olbermann and Sean Hannity. In this scenario, there would be no information sorting but high levels of information extreming.

Figure 2. *Conceiving of Information Sorting and Information Extreming as Two Distinct Dimensions*

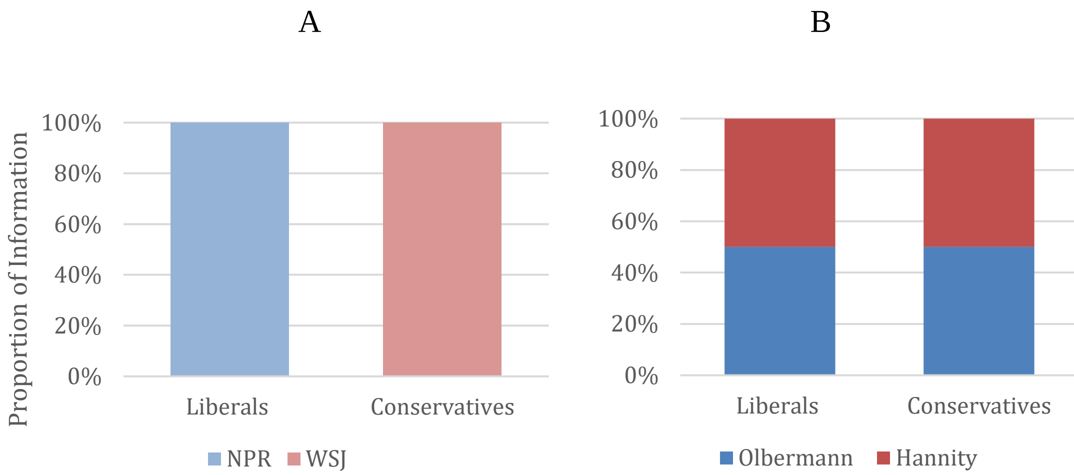

While information sorting and information extreming may go hand in hand, they are conceptually distinct, as illustrated by the two fictional worlds above. In fictional world A ideological users' information is completely sorted, with liberals only exposed to liberal information (from NPR) and conservatives only seeing conservative news (from the Wall Street Journal); but while information sorting is complete, the information users are exposed to is, by and large, not extreme. In world B, users are equally exposed to liberal and conservative information, thus with no sorting, yet the information they are exposed to is by and large extreme.

Extreming, like sorting, was not born of social media. Even during the Golden Age of "objective" media, news producers had incentives to report on potentially polarizing stories. News, by definition, is information that is out of the ordinary (Lippman, 1922; Shoemaker et al, 2001), and news gatekeepers have long known that the more extraordinary – or sensational – news is, the more it sells (Shoemaker et al, 2001; Soroka, 2012). In political coverage, the incentive to deliver sensational news often means reporting on "horse-race" politics, which not only highlights day-to-day political battles but also paints the parties as stark adversaries (Broh, 1980). In spite of those incentives to produce potentially extreme news, however, traditional media was also constrained - by long running journalistic norms, a few binding regulations and the power of advertisers - to present substantive, even-handed, and not too controversial coverage of political issues (Gans, 1979). Even though traditional media gatekeepers produced sensational news, they were limited in the extent to which they produced extreme news.

Those constraints largely fell by the wayside with the arrival of cable and the internet. With a fractured media and the disappearance of the Fairness Doctrine, as discussed earlier, news outlets could now diversify and present consumers with an array of viewpoints from across the political spectrum. Those views could be supported by substantive reporting, but more often than before "opinion reporting" – both easier to produce and juicier to consume – would flourish. The Glenn Becks, Keith Olbermanns, Drudge Reports and Daily Beasts that sprouted from the new media freedom (or free-for-all) created a breed of what the communications and political scholars discussed earlier have variously called "extreme," "partisan," or "outrage" media; media that aims to entertain while spurring emotions such as fear and anger, trading in "fervently populist or ideological rhetoric" (Prior, 2013), and liberally using "overgeneralizations, sensationalism, misleading or patently inaccurate information, ad hominem attacks, and partial truths about opponents" (Sobieraj et al, 2013; See also Levendusky, 2013; Taylor, 2017). In the 1970s it would have taken substantial effort to get access to such extreme news; even a subscription to Commentary or Motherjones would have been tame by comparison. But by the time social media arrived, however, extreme news content was free and plentiful (Baum & Groeling, 2008; Bennett & Iyengar, 2008).

While the more popular narrative is that social media cocoons us in political echo chambers, there is increasing recognition that social media sites are fertile ground for extreme information – and thus pose a threat to civil society. Speaking about social media in 2017, Barack Obama warned about "the capacity to disseminate misinformation, wild conspiracy theories, to paint the opposition in wildly negative light without any rebuttal—that has accelerated in ways that much more sharply polarize the electorate and make it very difficult to have a common conversation" (Remnick, 2016). Haidt (2016) similarly worries that "so long as we are all immersed in a constant stream of unbelievable outrages perpetrated by the other side, I don't see how we can ever trust each other and work together again." Journalists have made note as well. A Sunday NYT cover story opened its readers to the world of "hyperpartisan" "Facebook-native" sites that have come to dominate our feeds due to their adeptness at "cherry-picking and reconstituting the most effective tactics and tropes from activism, advocacy and journalism into a potent new mixture" (Herrman, 2016). So adept are they that Buzzfeed, an erstwhile media badboy turned responsible online journal, reported that those "sites now operate many of the partisan Facebook pages generating the highest engagement compared to more established players, and they therefore attract some of the biggest audiences online” (Silverman et al, 2016).

Extreme news has also caught public attention in its most virulent form as “fake” stories that often deliberately foment fear, anger, and inter-group conflict. Although not all or even most extreme news is “fake,” fake news often has the hallmarks of extreme news. The five most popular fake news stories leading up to the 2016 election were either endorsements of Trump (by the Pope) or allegations of illegal actions by Hillary Clinton (including murder and selling weapons to ISIS) (Silverman, 2016); each deeply tribal, emotionally evocative and by no means nuanced. The alarm caused by the rise of fake news stories in 2016, both their popularity as well as evidence that foreign actors had a hand in their promotion (Solon & Siddiqui, 2017), was arguably more audible than concerns over bubbles – and its response certainly more concrete. The social media giants, Facebook and Twitter, both took measures to weed fake news from their users’ feed, including hiring fact checkers and alerting users to less credible news sources (Constine, 2017). Yet nearly two years after the election US Congress, seemingly unimpressed with those efforts, still thought the dangers of fake news and foreign interference in social media warranted hauling Facebook and Twitter executives in to be grilled at Senate Intelligence Committee hearings (Romm & Timburg, 2018).

Outside the US, fake news that plays on emotions and inter-group distrust has likewise drawn attention as a result of its arguably more dire consequences. In India, fake stories on Whatsapp about kidnappers coming from other regions to steal children led to the murder of over twenty people whose only crime was being strangers passing through (Frayer, 2018). Viral Facebook stories in Sri Lanka likewise played on fears of others, but with a more direct out-group target; rumors of a Muslim plot to sterilize Buddhists there led to riots and mob violence (Taub & Fisher, 2018).

Scholarship, meanwhile, has only cursorily examined the rise and impact of extreme news in social media. Most academic research following the 2016 election has focused on quantifying and understanding the spread of fake news (Guess et al, 2018; Lazer et al, 2018). Yet some recent research suggests social media users tend to re-share information that is not necessarily fake, but is emotionally and morally charged (Brady et al, 2017). Users are both more likely to be exposed to stories about immoral acts on social media and more likely to have moral outrage triggered than if they were to encounter those stories offline (Crockett, 2017). We also know that Twitter users at the far ends of the political spectrum are often the most prolific posters of political information (Barbera, 2014). Studies by Facebook and Berkman researchers likewise confirm that news stories that come from media's extremes dominate what users share (Bakshy et al, 2015; Faris et al, 2017).[11]

But even more so than with information sorting, there is much left to understand about the potential of social media to amplify extreme views. Aside from circumstantial evidence just referenced, we do not yet know whether social media users do, indeed, prefer to share extreme as opposed moderate or balanced news stories. And if such a bias exists, current models do not tell us what effect diffusion will have on proliferating these extreme opinions.

Using the same approach as I do to examine information sorting, I investigate social media's capacity to promote extreme information. In contrast to Twitter users' absence of bias toward retweeting ideologically congruent information, I find a marked proclivity toward retweeting stories from extreme news sources. Using simulations on infinite and agent-based models I show that this bias will be amplified as information moves through homophilous networks and, unlike with information sorting, that there are no natural constraints on information extreming. Furthermore, while greater homophily has limited impact on information sorting, it does strongly amplify information extreming. In short, while the basic architecture of social media may be poorly designed to create ideological bubbles, it provides the mechanics for an efficient extreming machine.

[11] There is also a longstanding and rich scholarship on the prevalence of polarizing dialogue, particularly "incivility," on online platforms; we might guess that the forces that push social media users to engage in uncivil partisan flaming might also propel them to post extreme news (Anderson et al, 2013).

The Psychology of Sharing

In Part I of this dissertation, as just discussed, I break down the network mechanisms of social media to see the role each lever – diffusion, homophily and selective curation – plays in contributing to information sorting and extreming. In the second part, I zero in on one of those levers – selective curation – to better understand what drives social media users' selection of congruent and extreme news.

The question is not just academic. Feeling the heat from regulators and lawmakers, platforms such as Twitter and Facebook have begun to cull the worst forms of extreme news and to increase cross-ideological exposure. Facebook, for example, uses bots to either flag potential fake news and to censor or ban user accounts that promote hate speech (Constine, 2017). Twitter has also toyed with exposing users to a greater diversity of political viewpoints (Romm & Dwoskin, 2017).

But directly intervening to remove offensive or extreme speech is not the only approach platforms can take. Another is to nudge users to spread a more balanced and moderate information. To achieve balanced exposure, social media sites could encourage users to connect to other uses across the political spectrum, but doing so would not necessarily decrease the spread of extreme information. Influencing users' decisions on what to share and reshare, however, can be effective in both broadening users' information diet while reducing the intake of extreme information.

To influence behavior, we must first understand it. Remarkably, though, little is known about what motivates social media users to share political information online. Researchers have extensively studied the psychological drivers of sharing information more broadly (Berger, 2014), yet that research does not adequately explain the sharing of political information.

To fill that gap, I explore users' motivation for sharing political information on sites like Facebook and Twitter by developing and testing a theoretical framework to explain users' sharing behavior. That framework is built off the premise that, when posting political information on social media, users have an “imaginary audience” in mind, an audience that reflects one of their politically salient social groups, such as their fellow conservatives, environmentalists or Latinos. I propose two motivations that drive information sharing: “group impression management,” which signals users are true and loyal group members, or “group rallying,” which unites and incites their group in times of threat. Humans in general are motivated to project an image of themselves that helps them forward their social goals. On social media, when posting a political story, I posit that users’ goals are to secure inclusion and status within their politically salient groups. The image they aim to project to achieve those goals is that of a loyal and competent member. At times, however, users are also motivated to rally their political groups to address perceived threats, in particular those posed by the political outgroup. Both sets of motivations lead users to post stories with similar traits: at base a story will affirm the group’s beliefs and values, but those stories will also tend to be dogmatic in their position, emotionally evocative and tribal in nature, extolling the ingroup or denigrating the outgroup.

In a series of online studies, I test these propositions, specifically that users have their political groups in mind when they post political news on social media and that they do so to secure inclusion and status in those groups. I find some evidence consistent with the framework. Social media users do, indeed, tend to share political stories that affirm their group credentials – either because those stories affirm group values and beliefs, extol the ingroup, or denigrate the outgroup. I also find some evidence that users have a specific social group in mind as their imagined audience when they share news stories. I do not, however, find evidence that what specifically motivates subjects to share group-affirming stories is the drive to bolster inclusion in those groups.

Roadmap

The reader already has a good preview, from the discussion above, of what is to come. I take on the question of information polarization in social media in two parts. In Part I, I look at information polarization from the network-wide perspective and aim to answer the question “Is social media designed to increase information sorting or information extreming?” In Part II I focus on one of the key levers of sorting and extreming – selective curation – to better understand what drives users' tendency to share congruent and extreme information.

Part I begins with a discussion of the relationship between homophily, selective sharing and network diffusion dynamics and how they interact to determine information distribution patterns in social media. In Chapter 2, I present a series of diffusion models, using both mathematical analysis and agent-based modelling, to show what our expectations should be regarding social media's capacity to polarize our information environments. I find that, given different assumptions about homophily and selective sharing, we should have different expectations about a social media site's ability to increase information sorting and information extreming. I then examine what those assumptions should be, first by reviewing other scholars' research and analyzing public data to estimate levels of homophily on social networks on and offline in Chapter 3, and next by analyzing original data from the social media site Twitter to estimate users' selective sharing biases in Chapter 4. Finally, Chapter 5 inserts those estimates back into the diffusion models to answer the question "is social media designed to increase information sorting and extreming." I find that while social media creates the conditions to amplify extreme information, it is poorly designed to sort political information. Indeed, social media sites are more apt to expose us to a greater balance of political information than we might otherwise come across through other media although, again, it will be a balance of extremes.

In Part II, I switch gears from looking at macro patterns of diffusion to home in on the micro decisions of social media users. I also move from analysis of observational data and modeling to develop and experimentally test a set of psychological propositions. After ruling out existing theories to explain political information sharing, in Chapter 6 I develop an original theoretical framework to understand the motivations for sharing political information based on work in Sociology, Psychology, Anthropology and Communications. With a series of online experiments presented in Chapter 7, I then test elements of that theory.

The conclusion does what conclusions do: I aim to leave the reader with the importance of understanding social media's power to extremify our information environments and encourage other researchers to pick up where I have left off.

Part I: The Dynamics of Information Sorting and Extreming

Chapter 1. The Mechanics of Information Distribution

In the introduction I presented two narratives of social media and its capacity to polarize our information environments. In the more familiar narrative, social media acts as a sorting device, leaving users exposed to information that echoes back their views. In the less recognized narrative, users are not necessarily exposed to distinct and ideologically homogeneous political arguments as much as they are exposed to information that is “extreme”– dogmatic, emotionally charged and divisive - in nature.

How do we test these two theories of information polarization?

As previewed in the last chapter, rather than treating social media as a black box and comparing its outputs - in this case the extremity and ideological diversity of information consumers are exposed to - to those of other media systems, in this dissertation I look *inside* the social media box. There I aim to understand the factors - individual behaviors, structures and mechanisms - that might lead to information sorting or extreming.

There are any number of factors one could explore, but I focus on social media’s two defining features - users’ ability to connect to each other and their ability share information – along with a third feature that is a necessary by-product of the first two, the ability for information to travel from user to user. More specifically, I inspect users’ tendency to connect to others who share their political views (homophily), their ideological biases in what information they choose to share (selective curation), and the flow of politically valenced information through a homophilous network (diffusion) to see how those individual decisions and network-wide flows might produce information sorting and extreming.

The reader may wonder why it is worth bothering opening up the black box to answer the question “does social media contribute to information sorting and extreming?” when – as with Gentzkow and Shapiro (2010) – a more direct route would be to compare social media’s outputs to other media. I have three reasons which I explain below: 1) doing so can give us fundamental insights that outlive the current state of social media; 2) we may also be able to correct hidden – and misleading – intuitions we have about social media; and 3) most practically, we can give researchers and advocates a handle on how to correct the potential harms of social media.

The first reason is related to the features in the black box that I do not examine. While all social media sites are defined by the three dimensions above, each social media platform has a distinct set of features that also factor into how information is ultimately distributed throughout a network. In particular, many sites use algorithms to select which of a user’s friends’ posts are promoted in their feed. Some platforms likewise show ads and allow advertisers to target messages based on users’ suspected preferences. Both algorithms and targeted ads certainly have the potential to contribute to information bubbles and information extreming (Pariser, 2011), but those features are neither uniform nor stable; each social media company has its own algorithms for promoting posts and policies regarding ads, and those algorithms and policies can be changed on a dime.

The variability and changeability of platform policies creates two limitations. For one, it is difficult to claim, for example, that all social media platforms create information bubbles; there could exist sites that filter posts to insure its users are exposed to challenging information.[12] It would even be hard to claim that Social Media Site X, for example, fuels the spread of extreme information because while it may do so one day, it could implement a policy that downgrades all posts from far right and left media the next day, dramatically changing the ideological content of information on its platform. Facebook, for one, is known for making such feed-altering policies. In 2018, for example, the platform altered its algorithm to promote posts that produce more back

[12] While such social media platforms may not exist, there have been numerous attempts to build apps to help social media users break through their filter bubbles. (See, for some examples, Piore, 2018 and Lum, 2017.)

and forth engagement, which could have resulted in increased cross-partisan exposure (Mosseri, 2018). Another policy that year prioritized posts from more trustworthy sources, that likely decreased exposure to extreme information (Zuckerberg, 2018).

Examining the core mechanisms of social media, in contrast, gives us insights into what leads to information polarization both across platforms and over time, even as social media companies change their policies.[13] While platforms can easily change their algorithms and ad policies, they cannot change fundamental human behavior, such as what drives users to choose who they friend or what they post. They also cannot change the laws of network dynamics. By focusing on fundamental - and less changeable - features of social media I hope to gain insights that outlive social media's current state.[14]

The second reason to understand the inner mechanisms of social media is because we may think we already do so – and yet our intuitions may be leading us astray. Social networks are complex systems, and as such may run on rules that are not intuitive. We thus may have assumptions about the inner workings of social media that color our broader understanding of social media's impact. If inaccurate, those hidden assumptions could lay a false foundation for other theories about social media. But by dissecting the inner workings of social media we may simultaneously uncover what hidden assumptions we hold and dispel errant intuitions.

---

[13] It is true, however, that user behavior could change overtime as well; consumers may still be adapting to social media which is still less than a decade old (for most users), so their preferences about who to connect to and what to post may still be adapting. Yet those changes should be slow and will be constrained by fundamental human motivations.

[14] While I do not include an examination of algorithms to explain information polarization in this work, I by no means discount their role in contributing to or potentially mitigating information polarization. To the contrary, the choices platforms make on what information to promote or demote can substantially alter users' information environments; in the conclusion I will present some recommendations of my own on how to minimize information polarization. The relationship of the algorithms to the underlying mechanisms of social networks can be thought of as akin to the relationship between nature and nurture in behavior; the fundamental forces I look at in this work are social media sites' basic "nature," while the policies, design and algorithms platforms choose can "nurture" a range of outcomes.

To illustrate how complex systems can foil intuition – and how studying them can reveal and correct our hidden assumptions - consider Schelling's classic segregation model (Schelling, 1971). It is a well-known feature of US cities that they are racially segregated. Among the structural and social causes that may lead to high levels of segregation, one common intuition is that segregation occurs because people, particularly white people, prefer to live in neighborhoods where most others look like them. Schelling famously tested this intuition using an innovative method at that time – agent-based modelling - to see if, indeed, it was necessary for people to prefer segregation in order to create segregated neighborhoods. He set up linear neighborhoods in which he placed fictional residents ("agents") and simulated what would happen as those agents made decisions to stay on their plot or move to another house. Their decisions to move were based on whether a percentage of their neighbors shared their "race." If their neighborhood was above a given threshold they would stay; if not, they would move to the nearest suitable spot. His models showed, surprisingly, that in order for neighborhoods to be highly segregated, residents do not need to have a preference for segregation. To end up with neighborhoods that are 80% segregated, for example, it is only necessary that individuals have a preference to live in a neighborhood where at least 50% of their neighbors share their race. In other words, if everyone is happy to live in a mixed neighborhood (where up to 50% of their neighbors do not share their race), they will end up in highly segregated neighborhoods where only 20% do not look like them.

Schelling's model shows how the "micro" behaviors of individuals that interact in complex systems can lead to surprising "macro" patterns. As his segregation model illustrates, the "emergent" properties of a complex system can go against our intuition - or even reveal assumptions that seem so "true" we do not realize we had them. By examining the inner mechanics of social media, we may likewise reveal some of our hidden – and awry - assumptions and about how social media works.

The final reason for studying the mechanics of social media is the most practical. For bubble bursting advocates - or perhaps social media companies interested in mitigating extreming - it is useful to know how information sorting and extreming occur and what the strongest levers are to dial information polarization up or down. If we want to diversify users' information environments or decrease the spread of extreme news, does it make sense to focus on encouraging more cross-ideological connections or to instead nudge users' sharing behavior? Which effort will get better results? Understanding the underlying dynamics of information distribution will answer those questions and help practitioners direct their energies and strategies.

The Mechanisms of Social Media

So what is in the black box? And how does this dissertation attempt to shine a light inside? I've already identified the major parts and mechanisms are that I hope to illuminate - that is, users' choices about whom they connect to, their political biases in choosing what to post and, finally, how those two micro behaviors combine to create macro level patterns of information polarization. In the remainder of this introduction to Part I, I lay out a theoretical understanding how those three elements are related, how much scholars already know about each, and what is left for us to discover and understand. In the subsequent sections of Part I, I take up each of those elements and attempt to fill in the holes in our knowledge.

Homophily

Homophily, it is hard to dispute, is the linchpin of information sorting on social media; we place ourselves in information bubbles first – and possibly foremost - by surrounding ourselves with like-minded friends. Indeed, if we were to surround ourselves with diverse friends online, it would be difficult to conceive of how we would be exposed to homogenous information.[15]

The common and undisputed expectation is that social media users would, indeed, form homophilous online networks. It is well known that Americans tend to be friends with those who share their political views (Huckfeldt et al, 1995; Lazarsfeld et al, 1944; Lazer et al, 2010). This

[15] Though not impossible. We would have to assume either users have the capacity and willingness to selectively share one set of information with their liberal friends and another set with their conservative friends, or that social media providers insert a stringent ideology-sifting algorithm.

is not necessarily because individuals actively choose friends based on shared political outlook. Rather, secondary factors – shared community, same job, similar tastes in music, etc. – correlate with political beliefs and thus lead to friends sharing political outlooks (Huckfeldt & Sprague, 1987). We also inevitably influence – and are influenced by – our friends to think alike. On social media, to the extent that individuals connect to friends they have offline or seek out new like-minded friends we would expect to see users connect to others with similar political view (Benckler, 2006; Bimber, 1998).

Studies show that, indeed, our online selves are homophilic. Users on Twitter, for example, cluster by political views regardless of whether one considers networks by who retweets whom (Barbera et al, 2015; Conover et al, 2011; Grabowicz et al, 2012), hashtag use (Hanna et al, 2011), mentions (Conover et al, 2011; Grabowicz et al, 2012), key words (Himelboim et al, 2013) or who follows which political accounts (Boutyline & Willer, 2017). Facebook users who identify with a political ideology also, unsurprisingly, tend to friend co-ideologues (Bakshy et al, 2015).

There are, however, still open questions about homophily on social media. For one, it is not clear if our online networks are more homophilous than our offline connections. Theorists argue that social media give individuals both the capacity to find their political co-ideologues (Benckler, 2006; Bimber, 1998) as well as the serendipitous opportunity to encounter others they disagree with (Brundidge, 2010; Wojcieszak & Mutz, 2009). Whether those counter forces result in greater or less homophily is empirically unknown. It is also an open question how much homophily contributes to information sorting. As stated above, we can assume that homophily is necessary for information sorting to occur, but is it the dominant factor - or do sharing biases and diffusion play an equal or more important role?

Finally, while homophily can be expected to play a critical role in creating information bubbles, it is not known whether homophily would affect rates of extreming; if both your conservative and liberal friends have an equal bias toward sharing extreme information, will you be exposed to more extreme information when the composition of your friend network changes? We currently have no answer to that question.

Ideological and Extreme Curation

Just as it is hard to imagine information sorting without homophily, if users in an online social network didn't curate for extreme information, it is not easy to conceive how extreme information would gain a viral advantage. In order for one type of story rather than another to successfully proliferate through a network, users would need to have a preference for sharing that type of message.

Users would also need to have some bias toward sharing ideologically congruent information in order for information sorting to occur. Homophily on its own will not determine a network's level of information sorting. Indeed, homophily will only lead to information sorting if the politically aligned friends we connect to likewise post information that aligns with our shared views. If all your friends are conservative but they share equal parts liberal and conservative information, you will not be in a bubble. The diversity of information you are exposed to will thus be not only a function of how ideologically diverse your connections are, but the extent to which those connections curate information that aligns with their beliefs.

So far we have distinguished two types of curation - curating ideologically-congruent information and curating extreme information. Just as Sears and Freedman (1967) distinguished between "de facto" and "active" selective exposure, however, we also want to make a distinction between de facto and active curation. If, for example, a liberal social media user shares 70% liberal information, we would say she is a *de facto* curator, but we would not necessarily know if she *actively* selects information that aligns with her ideology. To be an active curator we would expect to see her re-post disproportionately more liberal information than she, herself, is exposed to - for example, if she were exposed to 70% liberal information but re-shared 80% liberal leaning news. Active curating, in contrast to de facto, demonstrates a bias in the selection of information to share.

While there is evidence that social media users are more often than not de facto curators of congruent information (An et al, 2014a; An et al, 2014b; Barbera et al, 2015; Shin & Thorsen, 2017), there is still no direct evidence that we actively curate ideologically-confirming information on social media. There likewise is scant empirical work to tell us if users have a bias toward sharing extreme information.[16]

Diffusion

The distinction between de facto and active curating matters when we consider the diffusion of information through a network. If we only looked at users' immediate connections and the information those friends share, we would be missing a signature feature of social media – its capacity to spread information in complex (and possibly "viral") ways. The information I receive in a network is not just determined by my friends and their sharing behavior, but also by my friends' friends and their sharing behavior, and their friends' friends, etc. We will want to know, then, what impact our extended network and diffusion will have on the information users are exposed to. Does diffusion through the network lead to higher or lower levels of information sorting and extreming? In order to answer that question, we need to know both the level of homophily of users in a network and users' active curation rate, but it is also necessary to model the flow of information in those networks to see what emergent outcomes arise.[17] Just as Schelling's model of segregation showed us that the micro behaviors of agents can lead to surprising macro outcomes, we might find that individual sharing behaviors have unexpected network-wide results.

Currently no such models exist to tell us the impact of diffusion on information sorting and extreming. Researchers have extensively modeled diffusion to answer other questions unrelated

[16] One exception being the aforementioned study by Brady and colleagues (2012) that shows that Twitter users do exhibit a bias toward re-posting tweets that are both emotionally and morally charged.

[17] In many ways, the potential mechanisms of information sorting on social media are nothing new; they are merely a re-rendering of incidental exposure and the "two step flow" of information first described by Lazarsfeld et al (1944) and later developed by Katz and Lazarsfeld (1955). As in Katz and Lazarsfeld's model, the information we are exposed to is a function of the people we surround ourselves with and the information they choose to share with us. What differentiates social media from the word-of-mouth media described by Katz and Lazarsfeld is that the information that flows now comes in the form of news links rather than news summaries (paraphrased by our friends). Social media also allows for more than a "two" step flow; with the ease of hitting a "retweet" button news can flow multiple steps. This ease might also diminish the role of "opinion leaders" who traditionally consume, synthesize and pass along information; now the less politically focused can get in the game of disseminating political information, for better or worse.

to the polarization of information, such as the relative role of strong and weak ties (Bakshy et al, 2012; Brown et al, 1987; Goldenberg, 2001; Grabowicz et al, 2012), the effect of hierarchical network structures on the size of cascades (Banos et al, 2012), and the kinds of beliefs that are more likely to diffuse via simple contagion, which requires one contact with that belief, or complex contagion, which requires multiple connections to have such a belief (Centola & Macy, 2007; Monsted et al, 2017; Romero et al, 2011).

There is likewise plentiful research that models other dimensions of network polarization, in particular studies that use agent-based models to explain how the influence of connections leads to the development of homophilous networks (Baldassarri & Bearman, 2007; Banisch et al, 2012; Bednar et al, 2010; Dandekar et al, 2013; Flache & Macy, 2011). Among this group, Sasahara et al (2019) illustrate how exposure to connections' posts on social media should lead to increasingly homophilous networks.

Two papers are worth noting for coming closest to answering our questions about diffusion in networks that have homophily or selective curation. Siegel (2013) uses models to demonstrate how diffusion can amplify media bias. In that work, however, bias is not ideological, but rather bias toward the non-status-quo position, and the networks modeled are not ideologically homophilous. Halberstam and Knight (2016) do model diffusion in a homophilous network in which nodes are ideological curators. They report several expectations, including that individuals will be exposed to more like-minded than challenging information and that information will reach like-minded users more quickly. Yet their paper likewise does not tell us how diffusion will affect levels of political information sorting and extreming in social networks.

Filling in the Gaps

In sum, while there is plentiful evidence that online social network users are homophilic and tend to share information that reflects their political viewpoints, there are still substantial gaps in our understanding of the mechanisms that might lead to echo chambers or information extreming on social media. To begin with, although we know online homophily exists, we do not know if levels of online homophily are comparable to homophily offline. Likewise, we know that ideologues tend to be de facto ideological curators, sharing information that mostly reflects their views, but we don't know if they are active curators of either ideologically congruent or extreme information. Finally, there do not exist models to tell us how diffusion in polarized networks - assuming some level of active curating - affects levels of information sorting and extreming.

In the next three chapters I attempt to fill in some of those gaps. I start by looking at networks globally, modeling diffusion to see if we should expect diffusion to amplify information sorting or extreming. I then review current research and take advantage of two public data sets to assess whether social media is evidently more or less homophilic than our offline networks. Finally, I analyze data collected from Twitter to look for evidence that social media users are active curators of ideological or extreme information. At the end of Part I, I pull together those findings to draw a first sketch of what happens inside the black box of social media and make predictions about social media's capacity to polarize our information environments.

Chapter 2. Diffusion

The previous chapter laid out the mechanics of information distribution in social networks and the three components we need to grasp in order to understand how information ultimately gets distributed. Those are: the levels of homophily in the network (how biased users are toward connecting to the like-minded); the curation behavior of users (users' biases toward sharing ideologically congruent and/or extreme views); and finally, the dynamics of diffusion (how users' sharing decisions build on each other to determine distribution patterns). Putting those three pieces together will give us a foundation for understanding how information is distributed across a social network - and what drives information sorting and extreming.

I start, in this section, by examining the last piece of the puzzle. I model diffusion on homophilous networks to answer several questions about diffusion's capacity to increase or decrease information sorting and extreming.[18] Assuming users are homophilic and have curation biases, are users' bubbles made more homogenous if they are placed in extended networks where information diffuses from connection to connection? Does diffusion lead to airtight chambers or are there limits to how much diffusion can filter our information and place us in ideological bubbles? Likewise, if users tend to re-share more extreme stories, how much - if at all - does diffusion favor the spread of extreme information over moderate information?

[18] I could have started instead by identifying and measuring levels of homophily and curation in existing social media, but before focusing in on specific online social networks, we may first want to get a broad picture of how homophily, curation and diffusion interplay. In doing so we not only improve our understanding of diffusion and are better able to predict levels of information sorting and extreming across different online social networks, but we may also find relevant patterns such as critical thresholds that tip the balance toward creating runaway information sorting or extreming.

As discussed above, existing models of diffusion cannot fully answer these questions. We also cannot trust our intuition for guidance. We may not be able to intuit, for example, how homogenous our information environments will be if 80% of our friends share our ideology and we all have a 90% bias toward re-sharing congruent information. Intuition likewise may not tell us how much extreme information will fill our feeds if people have a bias toward sharing extreme (as opposed to moderate) information. Given that online social networks are complex systems, the only thing we might expect is the unexpected.

To replace our intuition, I model diffusion on a range of homophilous networks. I first look at the dynamics of information sorting, starting with a simple infinite network that, although missing many of the features of real-life social networks, lays a foundation for understanding the effects of diffusion in polarized networks. I then move on to model diffusion on increasingly complex networks using agent-based models to see if the patterns of information sorting we see in simple infinite models hold across networks that more faithfully resemble real social networks. Finally, I look at information extreming, reversing course by first examining patterns in complex agent-based models and later infinite networks to understand why we might expect diffusion to lead to information extreming.

Information Sorting in Infinite Networks

To begin examining diffusion and information sorting, I use a simple network. As their name suggests, “infinite networks” have an infinite number of individuals (“nodes”), like the central node in Figure 3, that each links to a set of connections that branch infinitely in all directions. Although infinite networks lack many of the features we know exist in real-world social networks (for one, one’s friends cannot be connected to each other), I start here because they lend themselves to relatively simple mathematical interpretation. With these networks we only require algebra to gain insights into the fundamental dynamics of diffusion in homophilous networks. With such a grounding, we can then use those insights to grasp the dynamics in more complex and realistic networks.

I start with a few assumptions. First, information flows in one direction; each node has a set of connections from which it can receive information and another set to which it passes information along. Nodes are also either “conservative” or “liberal” and are homophilic, preferentially linking to co-ideologues. Finally, nodes are ideological curators; they prefer to pass along messages that align with their ideology.

Figure 3. *A Slice of an Infinite Directed Network.*

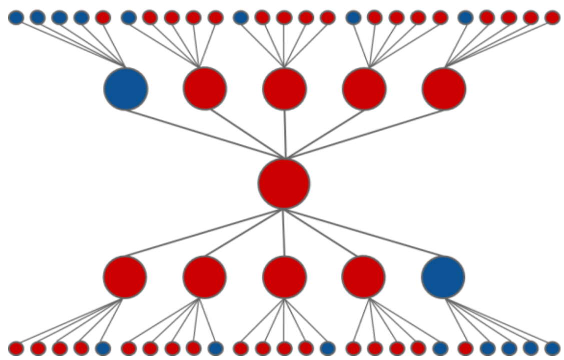

The infinite directed network we start with in this chapter is composed of an infinite set of nodes, each like the central node above. Each node is either conservative (red) or liberal (blue), receives messages from five friends from one direction and passes a subset of those messages along to five others. Both incoming and outgoing branches extend infinitely. 4/5ths of a node’s friends share its ideology. Note: Not all branches of the central node are depicted to avoid the image filling with an infinite number of nodes.

To see how information sorting may occur in this network, I simulate the diffusion of messages starting with one user posting either a liberal or conservative piece of information. In each simulation, once a user is thus “seeded” with a message, all its followers are “exposed” and each then decides whether or not to “re-post” (or be “infected” by) the message. Their decision to re-post is based on an ideological curating bias. If that bias is 90%, for example, a user would re-post messages that align with her ideology 90% of the time while posting counter-ideological messages only 10% of the time. If a user re-posts the message, their followers are in turn “exposed” and decide to repost or not – and the process continues.

What happens to levels of information sorting as a message diffuses through a network? First, let us mathematically define “information sorting” as the proportion of nodes exposed to a message that share that message’s ideology. Labelling the number of conservatives exposed as C_e and the number of exposed liberals as L_e, information sorting - when considering a conservative message - would be:

$$\frac{C_e}{C_e+L_e} \tag{1}$$

Simply put, if 70% of the nodes that are exposed to a conservative message are themselves conservative, then sorting is 70%. Later we will look at the flip side of sorting – that is, the proportion of messages any given node sees that aligns with her ideology.

Let us walk through one simulation, calculating the level of information sorting at each "wave" that a message moves through a network. In this simulation we use the following arbitrary parameters: a) 80% homophily in the network, b) an ideological curating rate of 90% for all nodes, and c) 50 friends for each node. To start, we seed a conservative node with a conservative message. In this initial wave, information sorting is 80%; that is, 80% of the nodes exposed to that post will be conservative and 20% liberal. This is a calculation we can do in our head, but since the math will soon get tricky, let us formalize the number of liberals and conservatives exposed to a message (L_e and C_e) as a function of how many liberals and conservatives have been infected by that message (L_i and C_i), how many friends those infected nodes have (*F*), and what proportion of those friends share their ideology (*H*) or do not (*1-H*):

$$L_e = L_i * H * F + C_i * (1-H) * F \tag{2a}$$

$$C_e = L_i * (1-H) * F + C_i * H * F \tag{2b}$$

In our starting wave, since C_i= 1 and L_i = 0, then L_e = 10 and C_e = 40. $\frac{C_e}{C_e+L_e}$, thus, is $\frac{40}{40+10}$ or 80%.

In the subsequent wave, the number of exposed nodes that go on to post the conservative message, and thus become infected, will be function of how many nodes are exposed and those nodes' curation bias: *S* (for "sharing") if the message shares a node's ideology and (*1-S*) if the message does not:

$$L_i = L_e * (1-S) \tag{3a}$$

$$C_i = C_e * S \tag{3b}$$

Plugging in a curation rate of 90%, that will mean 1 liberal and 36 conservatives re-post the message in the second wave and, using equations 2a and 2b again, that 400 liberals and 1450 conservatives are newly exposed to the message in this wave, making 78.4% of the exposed nodes conservative (and 21.6% liberal). In other words, information sorting has decreased from 80% to 78.4%.

We can iterate through this process a few times, but after just a handful of waves, we soon arrive at an equilibrium where 78.2% of exposed nodes on any given wave are conservative and 21.8% liberal. If we stick with an 80% homophilic network, but look at different rates of curation bias, we see in Figure 4a that an equilibrium appears to be reached no matter the rate of curation (or "selective reposting"). This is true regardless of starting point; if we instead seed a liberal node with a conservative message, as in Figure 4b, we see the same equilibria arrived at.

Figure 4. A Conservative Message Diffusing Toward a Sorting Equilibrium

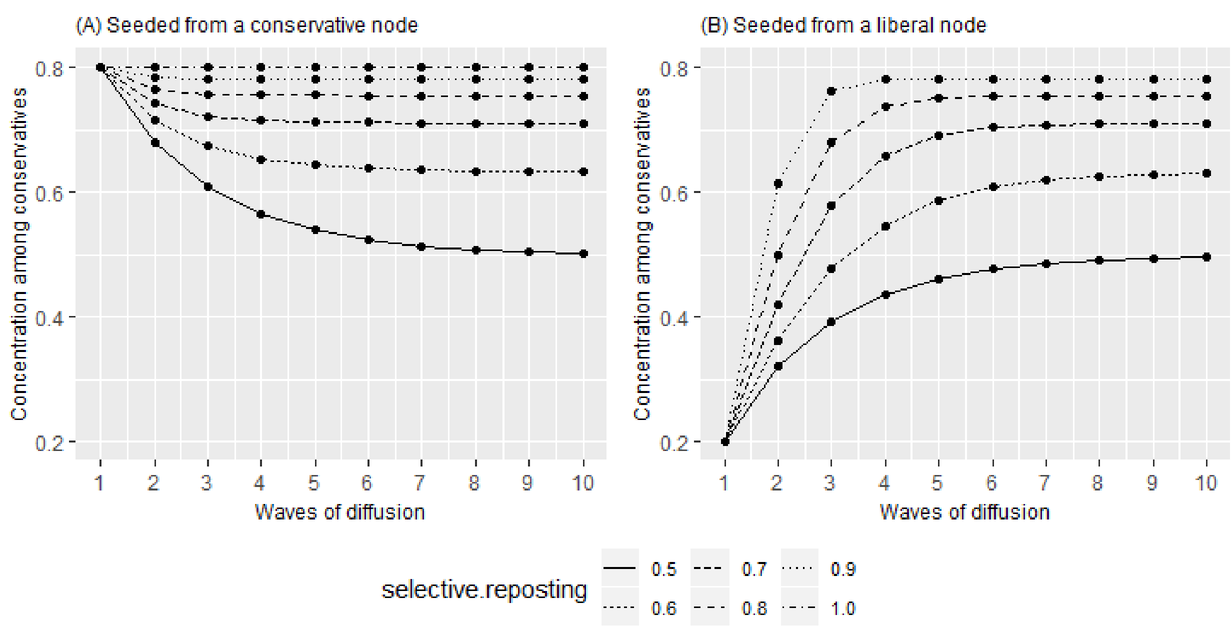

Information sorting levels ("Concentration among conservatives") of a conservative message as it diffuses through a homophilous network (of 0.8 homophily) after being seeded in a conservative node (A) or liberal node (B). Based on simulations in an infinite network. Each line represents diffusion for different levels of biased curation ("selective.reposting"), from 0.5 (no bias) to 1.0 (complete bias). Regardless of starting point, after a number of waves of diffusion, information sorting levels reach an equilibrium which is below the networks' level of homophily.

We arrived at these equilibria by iterating calculations for each wave. It is possible, however, to also derive the equilibria algebraically. To do so, we collapse formulae 2 and 3 to tell us the number of exposed nodes that are conservative or liberal in a given wave ($Ce_{t+1}\ and\ Le_{t+1}$) based on the number exposed in the previous wave ($Ce_t\ and\ Le_t$). Since we ultimately are interested in proportions we can drop the term for number of friends:

$$Ce_{t+1} = H * S * [\, Ce_t] \ + \ (1 - H) * (1 - S) * [\, Le_t] \tag{4a}$$

$$Le_{t+1} = (1 - H) * S * [\, Ce_t] \ + \ H * (1 - S) * [\, Le_t] \tag{4b}$$

At equilibrium, by definition, the proportion of conservative to liberal nodes that are exposed to a message will stay the same from wave to wave. To make it simpler to find that proportion – and to generalize the equation to not just refer to conservatives – substitute *P* (the proportion of nodes receiving a politically aligned message) for C_e, and *1-P* (the proportion of nodes receiving the same message, but for whom the message is counter-ideological) for L_e, giving the equations:

$$P_{t+1} = H * S * [\, P_t] \ + \ (1 - H) * (1 - S) * [\, 1 - P_t] \tag{5a}$$

$$1 - P_{t+1} = (1 - H) * S * [\, P_t] \ + \ H * (1 - S) * [\, 1 - P_t] \tag{5b}$$

Mathematically we know that:

$$\frac{P_{t+1}}{1-P_{t+1}} = \frac{H*S*[\, P_t] \ + (1-H)*(1-S)*[\, 1-P_t]}{(1-H)*S*[\, P_t] + \ H*(1-S)*[\, 1-P_t]} \tag{6}$$

Since again at equilibrium $P_{t+1} = P_t$, we can denote proportions at both time periods as *P*:

$$\frac{P}{1-P} = \frac{H*S*P \ + (1-H)*(1-S)*[\, 1-P]}{(1-H)*S*P + \ H*(1-S)*[\, 1-P]} \tag{7}$$

Finally, solving for P:

$$0 = (2S - 1) * P^2 + (2 - 2S - H) * P - (1 - H)(1 - S) \tag{8}$$

We can now plug in any H and S to the equation to calculate the level of information sorting at equilibrium. Figure 5 shows those equilibria across networks with different levels of homophily (lines) at different levels of ideological curation (dots on the line).

Figure 5. Sorting Equilibria across Networks of Different Levels of Homophily

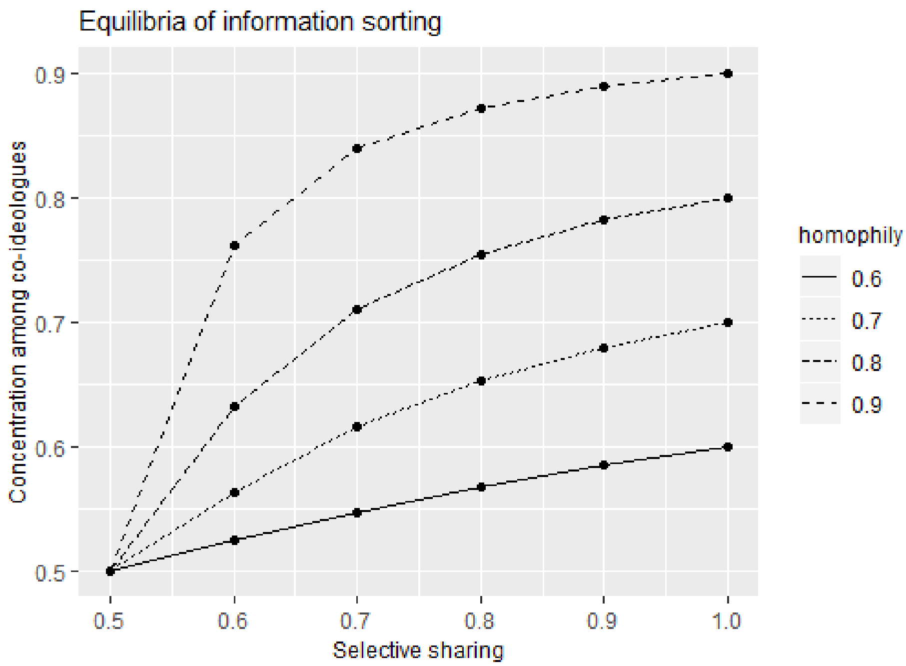

Each point represents the equilibria of information sorting ("Concentration among co-ideologues") reached in networks of varying levels of homophily (represented by lines) and biased curation ("selective sharing"). Equilibria are derived from simulations in an infinite network.

The first observation to make in looking at Figure 5 is that information sorting has limits; no matter the rate of active curation, information sorting equilibria never exceed the level of homophily in the network; equilibria will be somewhere between 50% (when ideological curation is essentially non-existent) and the level of homophily in the network (when curation is 100%, or absolute).

Figure 5 also shows that the relationship between curation biases and homophily is not linear; at high levels of curation bias, equilibria hew closely to homophily levels and do not alter much (i.e. have a shallow slope), but at low levels of bias equilibria diverge precipitously from the homophily ceiling. A takeaway for those interested in promoting diverse exposure to information is that altering levels of curation bias only has an impact when those biases are already relatively low; nudging users to go from 90% to 80% curation will not substantially diminish ideological information sorting whereas nudging users from 70% to 60% curation biases will have a larger impact in diversifying information exposure. At high levels of curation, then, a designer who wants to promote diverse exposure might do better to focus on promoting greater cross-ideological linking, whereas at low levels they will have more impact by encouraging less bias in curation.

So far we have observed, perhaps counterintuitively, that diffusion in homophilous networks does not result in runaway information sorting. Instead, diffusion leads to levels of sorting that are at least slightly more moderate than the level of network homophily. But the equilibria in Figure 5 only tell us the endpoint of diffusion. They do not tell us if diffusion *increases* information sorting. To answer that question, we would need to compare initial levels of information sorting to levels after diffusion. But what are those initial levels? In the diffusion simulations we calculated in Figure 4 we assumed that nodes initially either only posted messages that align with their ideology (which is somewhat unlikely) or only posted cross-ideological messages (entirely unlikely). In the first scenario, information sorting decreased with diffusion. In the second, information sorting increased. But surely individuals in social networks post initial messages that are somewhere in between being fully aligned with their ideology and fully incongruent.

A natural starting point may be to assume that users post initial messages with the same bias as they selectively re-post messages; it would be reasonable to assume that if users have an active ideological curation bias of 90%, their initial posts would also be 90% congruent with their ideology.[19] As seen in Figure 6, when we make that assumption looking at networks with 80% homophily, diffusion does increase information sorting, yet only marginally.

---

[19] In the next chapter we will see why this is not, indeed, a reasonable assumption. What it disregards is the set of information users are exposed to outside their network. Stay tuned.

Figure 6. *Diffusion's Effect on Information Sorting: The Starting Point Matters*

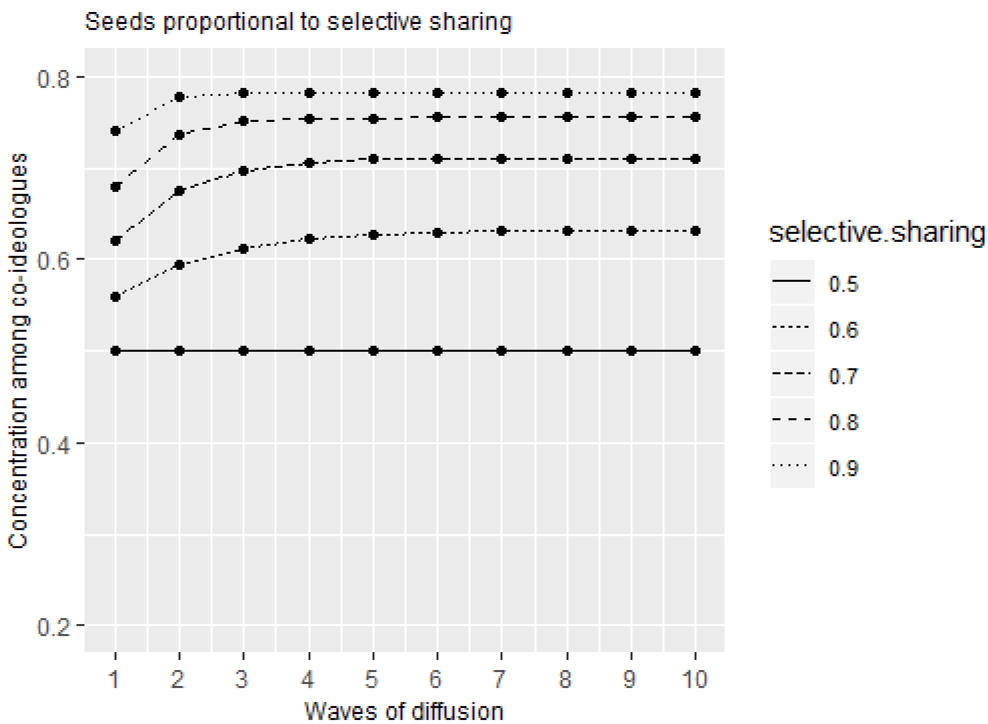

Information sorting ("concentration among co-ideologues") over waves of diffusion for different levels of curation ("selective sharing") in a network with 80% homophily. In this simulation, nodes are seeded with messages that match their ideology proportional to their selective sharing rates. The same equilibria are reached as in the simulations charted above, yet in this simulation equilibria are only marginally greater than levels of information sorting before diffusion.

To get a more complete picture of whether - and when - diffusion increases information sorting, we can plot the *change* in information sorting for any given initial posting bias and curation bias. Doing so for infinite networks with 80% homophily, Figure 7 reveals a few patterns. At the ends of the spectra, results are as we expect: if users have no initial bias in what they post but are 100% biased in what they choose to pass along, information sorting will increase by 0.3 (from 0.5, or no sorting, to 0.8, the level of homophily in the network). Flipping the biases around so users are highly biased in their initial posts but indifferent in their re-posts, sorting decreases by 0.3 (from 0.8 to 0.5).

Figure 7. *The Conditions in which Diffusion Leads to Information Sorting*

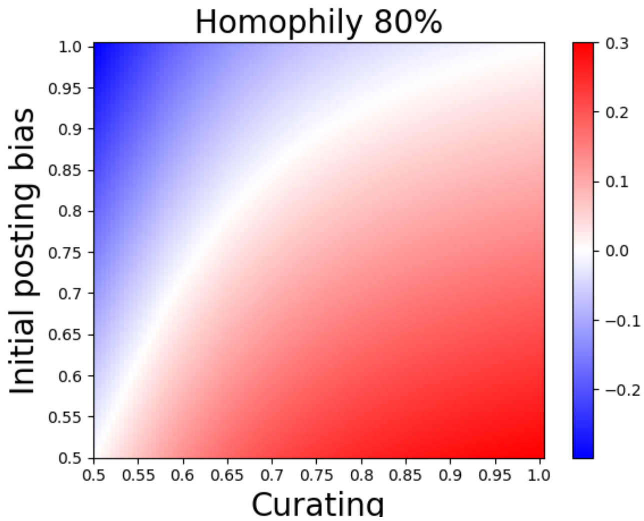

Whether diffusion leads to greater or less information sorting depends upon the selective resharing bias ("curating") of users and their initial posting biases. This chart shows under what conditions diffusion will lead to greater sorting (red) or less sorting (blue) in a network with 80% homophily. In this network sorting cannot increase by more than 0.3 (when curation is 1.0 and initial posting biases are 0.5) or decrease more than 0.3 (when curation is 0.5 and initial posting biases are 0.8).

A second observation is that there are more conditions that lead to increased than to decreased sorting levels. For diffusion to decrease levels of information sorting, curation biases must be lower than initial posting biases. How much lower depends on how homophilous the network is. Looking across networks of varying levels of homophily in Figure 8 we see that the greater the level of homophily the more often diffusion will lead to higher levels of information sorting.

Figure 8. *Conditions for Information Sorting across Networks with Varying Rates of Homophily*

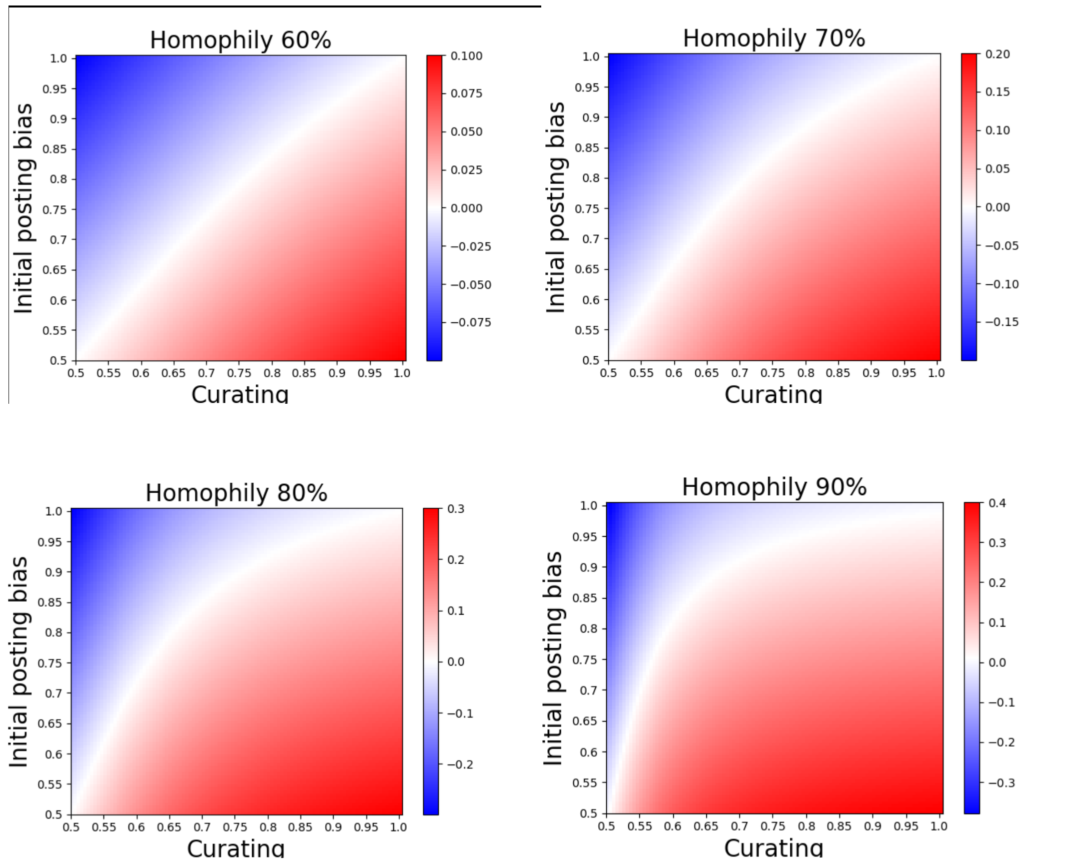

Just as in Figure 7, diffusion leads to greater or less information sorting depending upon the initial post bias of users and their curation biases. Looking across networks with different levels of homophily, however, we can see that diffusion leads to increased information sorting (red) under more conditions in networks with high levels of homophily.

These graphs, again, only show us the effects of diffusion we should expect given different assumptions about users' curating behavior and the network's level of homophily. To know if diffusion leads to greater information sorting in real social networks, we would have to know what those curation biases are and how homophilous the network is. We will identify those parameters in the next chapters.

Looking Up the Tree: How Diverse are the Messages we are Exposed to?

In the discussion so far, we examined information sorting from the perspective of the network as a whole, asking how evenly or skewed ideological information is distributed between left and right subgroups. But we may be more interested in looking at sorting from the perspective of the individual and asking if diffusion exposes users to more or less diverse information. Does the diversity of information individuals are exposed to likewise reach an equilibrium as messages diffuse through a network? Might we see airtight informational echo chambers form when looking from the individual's perspective? Or will we see a similar pattern as above with nodes being exposed to information that is more diverse than their friend network?

To answer those questions, we turn again to our infinite networks and algebra. Start by observing that one's exposure to ideologically aligned information will be a function of both homophily (H) and our friends' ideological curation biases (S):

$$Exposure\ to\ ideologically\ aligned\ information = H * S + (1 - H) * (1 - S) \qquad (9)$$

The first term $H * S$ represents the congruent information we receive from our co-ideologues and $(1 - H) * (1 - S)$ the congruent information we receive from our counter-ideological friends. But this formula does not give us a full account of the information we are exposed to. The S term here tells us how much congruent information our friends post *without regard to* how much ideologically aligned information they themselves receive; it assumes they receive equal parts liberal and conservative information. But since our friends have the same level of homophily as we do, we should assume they, like us, receive more congruent information from their friends. We thus need to calculate the proportion of congruent information users are exposed to given how much like-minded information their friends receive, which I will call P:

$$Congruence = \frac{H*S*P+(1-H)(1-S)(1-P)}{[H*S*P+(1-H)(1-S)(1-P)]+[H(1-S)(1-P)+(1-H)*S*P]} \qquad (10)$$

The top term represents the number of congruent messages we would receive; it is the same as the term in the previous equation, except adding in P, the proportion of congruent information our co-ideologues receive and (*1-P*), how much congruent (to us) information our counter-ideologue friends receive. The bottom term represents *all* the messages we would receive including congruent (same as the top term) and incongruent.

As we did earlier, we can derive the equilibrium using the expression above and noting that at equilibrium our exposure to congruent information will be the same as our co-ideological friends, or:

$$P = \frac{H*S*P+(1-H)(1-S)(1-P)}{[H*S*P+(1-H)(1-S)(1-P)]+[H(1-S)(1-P)+(1-H)*S*P]} \tag{11}$$

When we solve for P, we get:

$$0 = (2S-1)*P^2 + (2-2S-H)*P - (1-H)(1-S) \tag{12}$$

This is the same solution as previously derived when looking at information sorting from the perspective of the network. In other words, at equilibrium the level of information sorting in a network is the same as the degree of congruence for each node.

From the perspective of an individual, we can intuitively see why levels of information sorting will never be above levels of homophily in the network. Consider again a conservative node who has 80% conservative friends and 20% liberal; if we were to imagine that nodes were 100% biased in their curation – that is, that conservatives *only* forward conservative news and liberals liberal news – then that conservative node could not possibly receive more than 80% conservative information. The only way he could is if his liberal friends were less than 100% biased in the news they shared; but if that were the case then his liberal friends' friends would receive substantially more diverse information – and, network-wide, we would still end up with slightly less than 80% information sorting.

Information Sorting in Agent-Based models

The infinite models above are useful for giving us a basic understanding of diffusion dynamics in a homophilous network, but they are probably unrealistic for several reasons. All models are, by definition, simplifications of reality, but infinite models miss central features of real social networks that might make them particularly unrealistic and – perhaps - misleading. It may be that the simplicity of infinite networks is why diffusion does not lead to high rates of information sorting in those models. We might instead observe greater information sorting in networks that better represent the structure of actual social networks.

To see if it is merely the unique properties of infinite networks that are responsible for the results above, I simulate and measure diffusion on increasingly complex networks that more closely mirror real social networks. Unlike infinite networks, however, these networks are too complex for mathematical modeling. We cannot simply calculate the proportion of liberal and conservative nodes that will be exposed to a message as we did above; depending on which node is seeded and that node's unique set of neighbors and neighbors' neighbors, diffusion will result with a different distribution each time.

To simulate diffusion on these complex models I instead use agent-based models which allow one to construct complex networks and efficiently simulate diffusion hundreds (or thousands) of times. In doing so, we can measure the average distribution of information we should expect in different types of networks with different network parameters.

Agent-based models, true to their name, let us create sets of "agents" that can represent any kind of individual entity (such as an individual, a company, or an ant). Those agents are able to take specific sets of actions - move residences, buy parts, move left or right to find a grain of rice - based on individual decision algorithms. What makes agent-based models useful – and interesting – is that agents can be placed in environments with other agents and base their decisions on the actions of those agents. Simulations can thus reveal emergent macro behaviors that result from the micro decisions of a population of agents interacting with each other. In the simulations in this chapter, agents take on the role of social media users, their actions are the ability to "post" political information, their environment is a network of connections and their decisions are whether or not to post a message when one of their friends does so.

Agent-based models could be configured in any programming language. For the simulations described below I use Python which, among other attractive features, includes the module Networkx that allows for the easy manipulation of networks.

Types of Networks

I simulate diffusion on three types of networks ("graphs" in network science speak), each with increasing complexity and additional attributes to make them more closely resemble real social networks. As with the infinite networks discussed above, all graphs are constructed with varying levels of average homophily and with nodes that have active ideological curation biases; however, in these networks a graph's level of homophily is an average among all the nodes. Also unlike the infinite networks we saw, these graphs are "undirected"; that is, messages can flow in either direction between two linked nodes. The three types of graphs - and how they are constructed - are described below.

Binary Random Graphs.

In random networks, also known as Erdos-Renyi graphs, all pairs of nodes are linked with a given probability (Erdos & Renyi, 1960). Random graphs move us closer to real social networks in a few key ways. Unlike in infinite networks, it is possible for friends of a given node to be friends with each other. Random graphs also exhibit a signature feature of real-life social networks; to move between any two nodes in the network, a message only has to hop through a small number of nodes. In an infinite network, it would take an infinite number of hops. You may be familiar with this characteristic of social networks as the "six degrees of separation" that can make even the largest network feel like a "small world."

To create a binary *homophilous* random network, I vary the Erdos-Renyi graph by first assigning nodes a binary ideology of either "liberal" or "conservative" and then adding a homophily bias into the probability that any two nodes link. For example, if the homophily of the network is 80% then the likelihood that two nodes that share an ideology will form a link is $0.8 * CP$ (where CP is some ceiling probability) while the likelihood that nodes which do not share an ideology will be linked is $(1 - 0.8) * CP$. Figure 9 shows such a graph: one can see how like nodes appear to group with like.

Figure 9*. Example Random Binary Network*

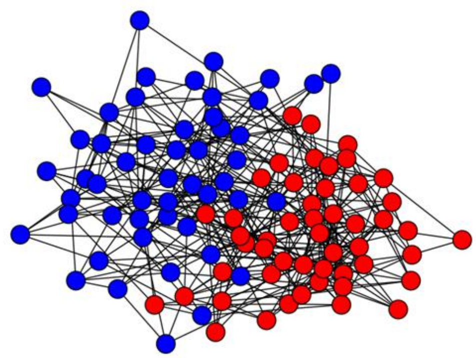

Example of a binary random network with roughly equal numbers of conservative (red) and liberal (blue) nodes.

Binary Small World Graphs.

Random graphs, while more realistic than infinite graphs, still lack important characteristics of social networks. For one, in social networks we know that friends of friends not only *can* be friends with each other, but are *likely* to be. In graph theory this phenomenon is called clustering or "triadic closure" (when your two friends become friends, the triangle is closed). Random graphs, in contrast to real world social networks, have little clustering. Another way in which random graphs are dissimilar to social networks is that all nodes have similar levels of popularity; some may have more friends than others, but there is little variance and that variance hovers around an average (statistically speaking, popularity is normally distributed around a mean). In social networks, however, we know that the distribution of friends is skewed; that is, a few nodes tend to be very popular while most nodes are relatively unpopular.

To generate graphs that include high clustering and a skewed friend distribution, I apply two common methods. First, I construct a Watts-Strogatz "small-world" network (Watts & Strogatz, 1998). These networks are built by starting with a ring of nodes that are all connected to their nearest neighbors, thus ensuring that there is a high level of clustering (since a node's neighbors will also be connected to each other). But such a ring is not yet small world since, unlike the random graphs above, it would take many hops for a message to move between a random pair of nodes. To add the "small worldness" to the highly clustered ring, Watts and Strogatz rewire a subset of the connections, randomly breaking links between neighbors and re-connecting nodes to a random node in the network.

These small world networks get us two qualities of social networks - high clustering and relatively few degrees of separation - but nodes still have about the same numbers of friends. To create skewness in the distribution of numbers of friends, I add an element of "preferential attachment" into the rewiring process. In preferential attachment, when nodes link to each other they prefer to connect to nodes that are already more popular. If nodes are rewired sequentially, popular nodes become increasingly more popular with each new connection.

Finally, to bake in ideological homophily, I generate the initial string of nodes using a given probability that each successive node shares the ideology of the previous node; by trial and error I determine what that probability needs to be in order for neighbors to have a given average homophily. Likewise, in the second step of Watts-Strogatz graph generation, nodes are re-wired to nodes with a probability biased toward connecting with co-ideologues.

In Figure 10, you can see the result: as with the random graph above like group with like, but now you can see triads forming and heavier clustering. There also appears to be greater variance in the number of friends each node has (an attribute known as "node degree") than in random graphs; while the graph below has the same average node degree as in Figure 9, in the small world network there are nodes with only 1 or 2 friends.

__Figure 10__. Example Binary Small World Network

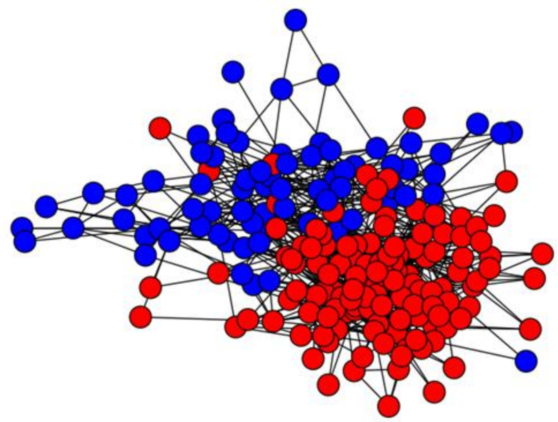

Example of a binary small-world network with roughly equal number of conservative (red) and liberal (blue) nodes.

Continuous Random Graphs.

Finally, because individuals' ideologies are not binary in the real world, I create graphs in which nodes take on an ideological score anywhere between 0 and 1. As in the binary random graphs, in these graphs two nodes are linked with a ceiling probability multiplied by a homophily bias. In constructing these graphs, that homophily bias is based on how ideologically distant the two nodes are from each other; specifically, the bias is set so that a node's friends will be normally distributed around its ideal point. To vary the level of homophily in a given network I altered the standard deviation of those individual friend distributions; for example, to create a graph where 90% of nodes' friends shared their ideology, I used a standard deviation of 0.17 while for a network with 70% average homophily a standard deviation of 0.48. (Those standard deviations were "calculated" via manual trial and error.) An example of such a graph is seen in Figure 11.

Figure 11. *Example Continuous Random Graph*

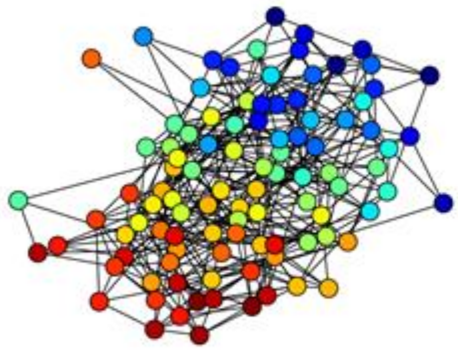

Example of a continuous random graph with nodes of ideologies from extremely conservative (dark red) to extremely liberal (dark blue).

Simulating Diffusion

For each network type described above, I simulate diffusion on networks with varying levels of homophily and ideological curation biases, just as we did with the infinite graphs. Again, because diffusion is a complex process, each simulation will result in a unique distribution which depends not just on which node is seeded but also the unique network constructed. To find the expected – or average – level of information sorting for each set of network parameters (type of network, homophily level and curation bias), I generate 100 graphs and run 200 simulations, giving a total of 20,000 observations. All graphs have 1,000 nodes with an average degree of eight links per node.[20]

Each simulation starts with seeding one message (tagged as conservative or liberal) in one node. As with the infinite network model, that node's connections are "exposed" to that message and then use a decision algorithm to decide whether to re-post the message. For the binary graphs, that decision algorithm is a ceiling probability (of 10%) times *S*, the curation bias, if the message

[20] I run the simulations with average degree of 6 and 10 as well and find similar results as the ones reported here.

is congruent, or (*1-S*) if the message is incongruent.[21] I run simulations with curation biases of 0.5, 0.65, 0.8 and 0.95. In the continuous graph, those curation biases are applied to messages that are closer to the node's own ideology; I use a normal probability distribution so that messages closest to the ideology of the node have a high probability of being re-shared by that node. I vary the degree of "closeness" by using a range of standard deviations of 0.1, 0.2, 0.3 and 0.4.[22]

At each wave of diffusion, I capture the number of liberal and conservative nodes that are exposed to the message. Each simulation runs for five waves.[23] At the end of the 200 simulations for each graph constructed, I also calculate the average diversity of information that nodes are exposed to at each wave.

A Note on Simple and Complex Contagions

In the diffusion models used in this section, a node can have multiple opportunities to become infected; each time a new friend posts a given political message it can decide anew whether or not to re-post that message (although once a node posts a message itself it cannot become infected again).

This type of contagion differs from "simple" contagion models in which nodes can only become infected at first exposure. Such simple contagion models are used, for example, in understanding the spread of diseases in which individuals, when first exposed to a virus, are either infected or proven to be immune. The spread of information can likewise follow a simple contagion pattern;

---

[21] I chose this ceiling probability for two reasons: it was low enough to ensure that networks would not be easily saturated yet high enough to ensure messages would not easily die out. In Chapter 5, I use a ceiling rate that more accurately reflects real re-sharing rates.

[22] I also run simulations in which the decision algorithm is based off a fixed threshold distance or decreasing linear probability. I show results from the normal probability distribution as they exhibit the most extreme bias and so – I hoped - were more likely to produce sorting results.

[23] I choose five waves mostly for reasons of coding ease, but it also is an appropriate number; we want to see what happens over time, yet even the most successful viral memes rarely make more than five hops (if only because individuals are all connected, once again, by "six degrees of separation").

if there is a type of news story that a social media user knows instantly is relevant to their network then they will either share it the first time they observe it or never at all. Such messages might be breaking news such as the death of a loved celebrity or the results of an awaited political outcome.

The spread of political news, however, might not always resemble a simple contagion, but instead diffuse as a “complex” contagions in which individuals must be exposed multiple times to a news story before sharing it. Research suggests that behaviors often spread as such complex contagions (Centola & Macy, 2007). Taking on a new behavior can be risky; the new movie you go to see may be a waste of your time, or the new hair style you try may be embarrassingly unhip. But the more friends who have gone to that movie (and perhaps have said they liked it) or the more that sport that hairstyle, the less risky that behavior is to adopt. Posting a political message likewise entails risk; because political information, particularly information about new stories and issues, may be difficult to parse or put in context, it might not be as clear that a bit of news is worth sharing. Posting an irrelevant or inappropriate story may be a waste of energy or make one look foolish. If, however, one sees a story posted multiple times by friends, the significance of that story will become more certain.

Which kind of contagion should be modeled for the spread of political information? There is evidence that information diffusion on social media conforms more to a complex contagion model (Monsted et al, 2017), however that may only be the case for certain types of information (Romero et al, 2011).

In the simulations in this section, I use a diffusion model that splits the difference between simple and complex contagions. Most models of complex contagion require a node be exposed to a threshold number (or percentage) of infected connections before becoming infected itself. In the simulations I use I allow nodes to be infected with a single exposure to a news story (a la simple contagion), yet if a node is not infected on first exposure it has a new chance to become infected each time a new friend shares a given news story. This way I allow for simple contagion yet at the same time account for the likelihood that social media users are more apt to share a political story the more they are exposed to it.

Results

Running simulations on networks of varying degrees of homophily and seeding nodes with messages that align with their ideology, we can compare information sorting trends to those we saw in infinite networks. As we can see in Figure 12, diffusion leads to decreased information sorting at each wave (again, when seeding a node with a congruent message) and appears to approach the same equilibria we saw with the infinite networks. There is one exception; in continuous networks, when nodes have a strong ideological curation bias (with a bias toward messages that are no farther than 0.1 ideological points away), diffusion sometimes begins to tip toward sorting levels greater than the network's level of homophily.[24] Yet, as the simulations only run for five waves we cannot see if this trend will continue and break through the homophily ceiling.

[24] I show only the results of simulations run on networks with homophily levels of 0.8 to be consistent with the infinite models and not clutter these pages. I, however, run simulations in networks with other levels of homophily and see similar results, which can be seen in Appendix I.

__Figure 12__. Diffusion in Agent-Based Models: Information Sorting at the Network Level

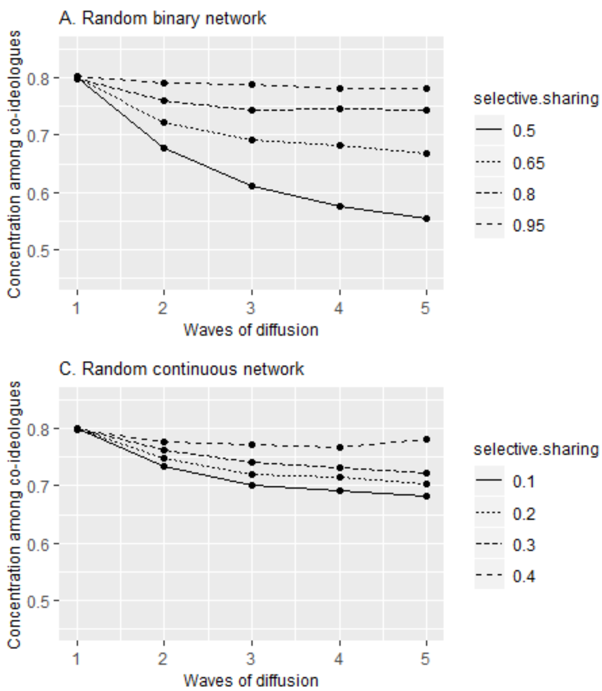

Information sorting ("concentration among co-ideologues") over waves of diffusion for different levels of curation ("selective sharing"), in simulated diffusions on a) binary random networks, b) binary small world networks and c) continuous random networks. In these simulations, just as those modeled on infinite networks in Figure 4a, nodes are seeded with messages that share their ideology. Also as in Figure 4a, all networks have average homophily of 0.8 and average degree of 8.

If nodes post initial tweets with the same bias with which they re-post messages, we again see (in Figure 13) similar sorting patterns as with the infinite network models: information sorting increases slightly and appears to approach an equilibrium. (I do not re-run the models on continuous networks as they only post the initial tweet with their re-sharing bias.)

Figure 13*. Diffusion in Agent-Based Models: A Selective Starting Point*

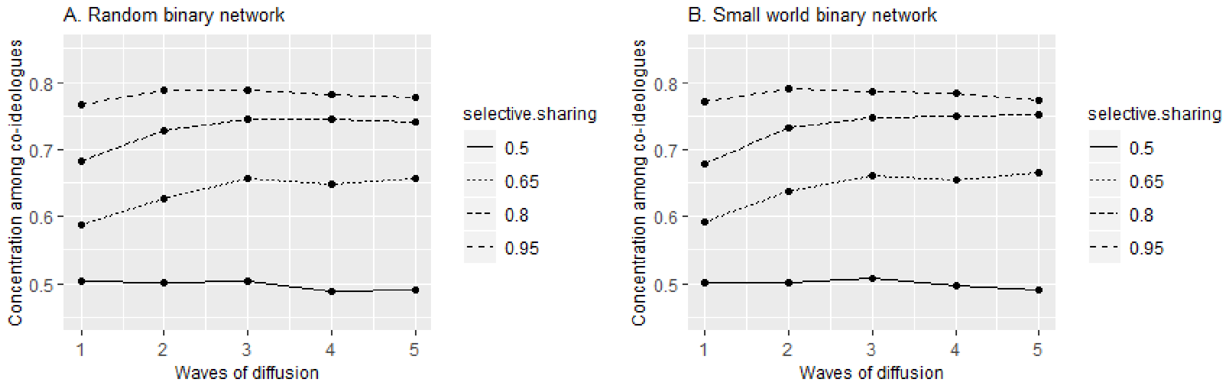

Information sorting ("concentration among co-ideologues") over waves of diffusion for different levels of curation ("selective sharing"), in simulated diffusions on a) binary random networks and b) binary small world networks. In these simulations, as opposed to those in Figure 12, initial messages share the ideology of seeded nodes relative to the nodes' curation rate. As with the models in Figure 12, all networks have average homophily of 0.8 and average degree of 8.

Finally, the agent-based models let us observe information sorting from the perspective of the individual nodes. By tracking the diversity of messages nodes are exposed to at each wave of diffusion, we can see if diffusion results in nodes being exposed to more or less diverse information. Figure 14 charts the diversity of information nodes are exposed to at each wave of diffusion. As predicted by the mathematical equilibria in our infinite networks, information sorting at the individual node level mirrors sorting patterns at the network level.

Figure 14. *Diffusion in Agent-Based Models: Information Sorting at the Node Level*

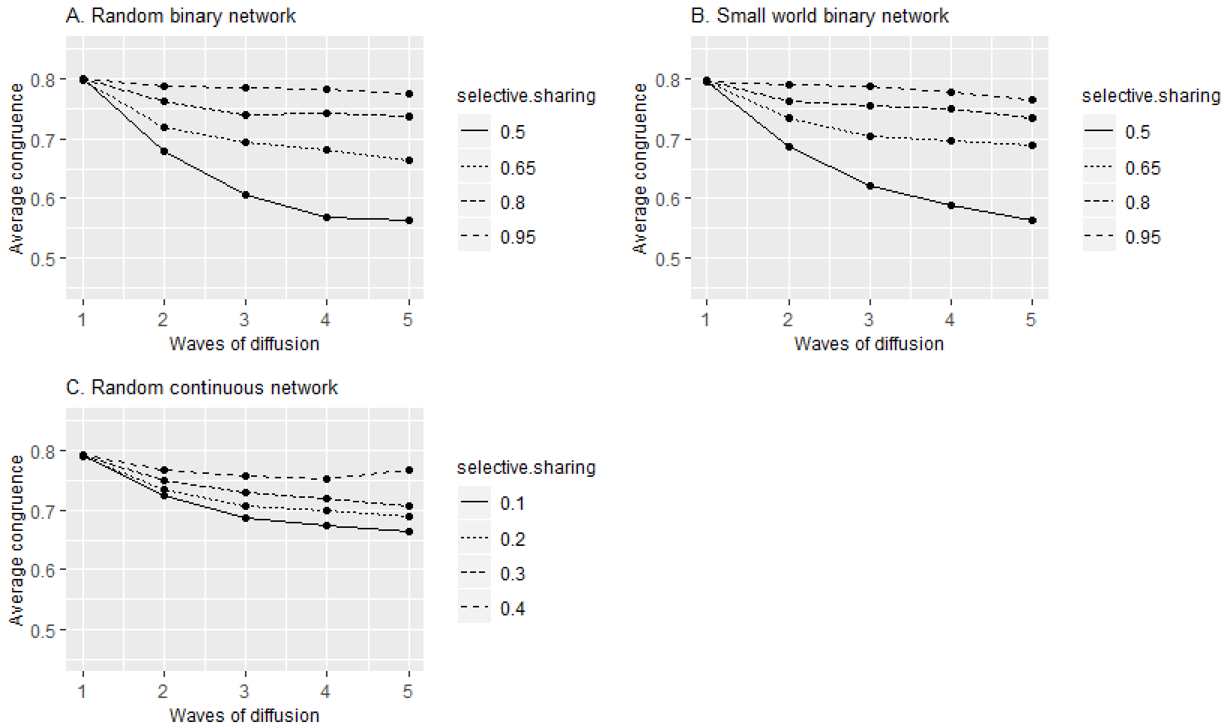

Proportion of messages nodes are exposed to that are ideologically congruent at each wave of diffusion for different levels of curation ("selective sharing"), in simulated diffusions on a) binary random networks, b) binary small world networks and c) random continuous networks. Initial messages share ideology of seeded node. Figure 12 looks at information sorting at the network level, i.e. how much ideological messages are concentrated among users who share the messages ideology. In these graphs, we instead look at information sorting at the level of the nodes; at each wave of diffusion the graphs show the proportion of messages a node will see that shares that node's ideology.

The agent-based models and simulations above suggest that the results we saw in the infinite network models are not anomalies but may rather be the norm. These are only three network types but, with the possible exception of one set of parameters on one network type, simulations all point to the observation that information sorting, again, cannot exceed the level of homophily in the network.

One cannot, of course, say that diffusion leads to levels of sorting lower than a network's homophily in *all possible* networks; to do so we would need to construct an infinite variety of networks on which to run simulations – and still there could be the black swan network we missed. If such a network model exists, however, it would only be relevant if its structure resembles real social networks. But for now, given the models above and without evidence to the contrary, it is reasonable to assume that diffusion has a limit – i.e. the level of homophily in a network – to how much is can sort left information from right.

Information Extreming in Agent-Based Models

Diffusion's capacity to strengthen information bubbles, as we have just seen, seems minimal at best. Even at the highest levels of filtering, models show that diffusion does not lead to airtight echo chambers; rather, information sorting as a rule will be capped at a network's level of homophily.

But while diffusion has limited capacity to increase information sorting, what of information extreming? How much - if at all - does diffusion favor the spread of extreme information? And is there a ceiling to how much extreme information can be amplified, as we saw with information sorting?

Before answering those questions, let us revisit the distinction between information sorting and information extreming. I defined information sorting as the proportion of users exposed to a message who share that message's ideology or, from the individual's perspective, the homogeneity of messages individual users are exposed to. In other words, the more information sorting there is the purer the information bubbles. Information extreming, in contrast, is concerned with the proportion of *extreme* messages to *moderate* messages that flow through the network and that individuals are exposed to. When levels of information extreming are high nodes will be exposed to high proportions of extreme information - from either side of the spectrum.

By extreme information, recall, I do not mean information that reflects *ideological* extremes, such as communism or fascism; rather I define extreme information as stories, articles and memes that support liberal or conservative points of view in a way that is dogmatic, tribal and emotionally-valenced.

To see how diffusion might amplify extreme information, we can build off our information sorting models in binary networks, but with a modification. In those models, we only considered binary messages – i.e. messages that are either conservative or liberal. We assumed that users have a simple inverse bias in sharing ideological information - so if a user has an 80% curation bias she is roughly four times as likely to share an ideologically congruent than incongruent story that she sees (sharing an aligned story 80% of the time and an unaligned story 20% of the time). By "ideologically congruent," I merely mean any message that has a left bent (for a liberal) or right lean (for a conservative), regardless of how moderate or extreme that message is. We can visualize agents' sharing decisions in those models in Figure 15, again using a network where agents have an ideological curation bias of 0.8: if a message is anywhere along the liberal continuum (< 0.5 on an ideological scale) liberal users will share that message 80% of the time while sharing conservative messages (> 0.5) 20% of the time – and conservatives vice versa.

***Figure 15**. Conceptualizing Binary Selective Curation*

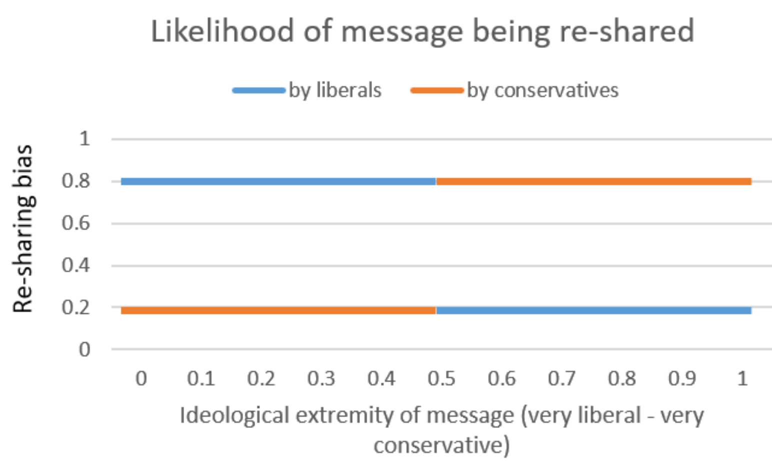

This image conceptualizes one way of modeling curation biases of liberals and conservatives across messages of ideological slants from very liberal to very conservative. In this model, ideologues have binary biases. Liberals prefer to share liberal messages 80% of the time and conservative messages 20% of the time. That preference is uniform for all liberal and conservative messages; it does not matter if a message is moderate or extreme.

To simulate the diffusion of extreme messages, however, we need a way to differentiate not only between liberal and conservative messages but also between moderate and extreme messages. It is also necessary to model how users might decide to share extreme versus moderate messages.

There are at least two ways to do so. One would be to assume that individuals have a bias toward re-sharing *any* extreme information, regardless if it comes from the left or right. But this assumption would be both trivial and unrealistic. Trivial because if, for example, we were to say that all users in a network are twice as likely to re-share a Breitbart story over a Newsweek story, we know already from existing growth models that the Breitbart post will spread exponentially faster than the Newsweek post and quickly dwarf the reach of the Newsweek story.[25] It would also be a stretch to assume that all individuals would be biased toward sharing a post from Breitbart or, to pick an example from the left, DailyKos - when, surely, conservatives are more likely to prefer to share the Breitbart and liberals to share the DailyKos story.

More realistic is to imagine instead that a conservative user, for example, has sharing preferences that vary across the spectrum from extreme right to extreme left. He might have the strongest bias for sharing a post from Fox News, which strongly confirms his worldview, and a slightly weaker bias toward sharing a story from a moderate conservative publication like Forbes which would be less dogmatic in its support of a conservative worldview. Moving leftward, that conservative user would be even less apt to post an article from a moderate liberal publication like Newsweek and, finally, would not be caught dead posting an article from Huffington Post (though they might do so to mock the story - but that is another matter). Such a bias is schematically visualized in Figure 16.

[25] Assuming the resharing rate is high enough that both stories don't both die out.

Figure 16. *Conceptualizing Extreme Selective Curation*

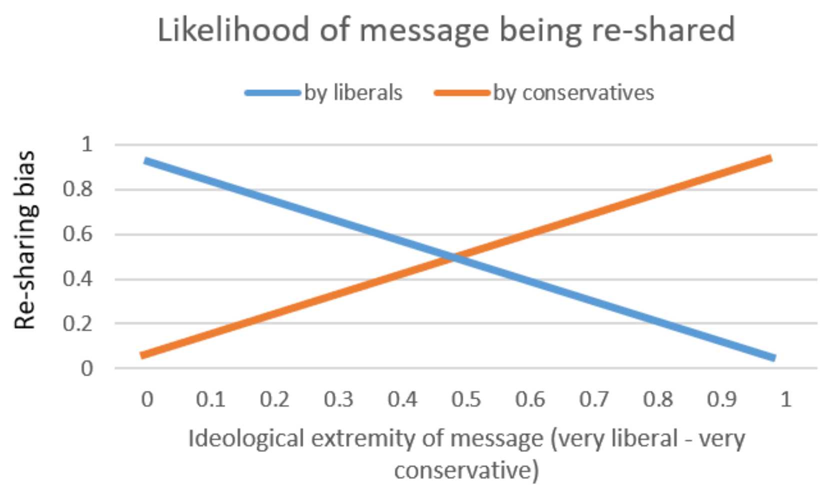

A second way of modeling curation biases of liberals and conservatives across messages of ideological slants from extreme liberal to extreme conservative. In this model, ideologues' biases are both ideological and relative to how moderate to extreme the message is.

While more realistic, this model has its own simplifications that need be acknowledged. First, it assumes that the decision to share a story is merely dependent upon the ideological extremity of the story. There are, of course, countless other criteria individuals may use to decide to re-share a story - a post could be funny, awe-inspiring or surprising, it might contain a photo, etc. In this model we assume all those other qualities even out. In other words, "all else being equal" users will be biased towards sharing stories that more strongly confirm their ideological and partisan views.

Second, while the reader might agree in the general direction and angle of the lines presented in Figure 16, they might suggest minor alterations. It might be, for example, that there is a threshold beyond which the extremity of a message becomes off-putting; while a conservative might happily share a story that slams Obama's accomplishments in office, if the story adds a racial layer it might make them pause. Similarly, a liberal might be apt to share a meme exhorting friends to fight climate change, but would hold back if the meme recommended using violence. We could also imagine individual users' thresholds varying. But while these modifications – and others I have not imagined – are certainly plausible, once again, I use this simplified schema for the sake of parsimony.

With this decision model, we can now compare the diffusion of extreme and moderate information using the same simulation models above. Now, instead of assuming nodes have the same ideological sharing bias for all message left or right of neutral (as in Figure 15), in these simulations nodes' curation bias is relative to the extremity of a message (as in Figure 16). Also, instead of looking at the relative number of conservatives and liberals exposed to a message, we compare the number of all nodes (conservative and liberal) who are exposed to an extreme vs. moderate message over successive waves of diffusion.

I run simulations that diffuse messages of varying extremity from neutral (0.5) to extreme (0.95) on networks of varying levels of homophily. The results for networks with 80% and 95% homophily are shown in Figure 17. In a network with 80% homophily, in Figure 17a, we see a strong relationship between the ideological extremity of the message and the number of exposed nodes. Neutral messages (0.5) almost completely die out after 5 waves, while extreme messages (0.95) continue to proliferate (though at a decreasing rate). By the fifth wave, approximately five times as many nodes are exposed to extreme messages as they are to neutral messages.

If we consider a network with higher levels of homophily, such as 95% in Figure 17b, we see that the gap between exposure to extreme as compared to neutral messages widen further, with extreme messages outpacing neutral message by a ratio of almost ten to one by the fifth wave.

Figure 17. *The Viral Edge of Extreme Messages*

A. Network of 80% Homophily

B. Network of 95% Homophily

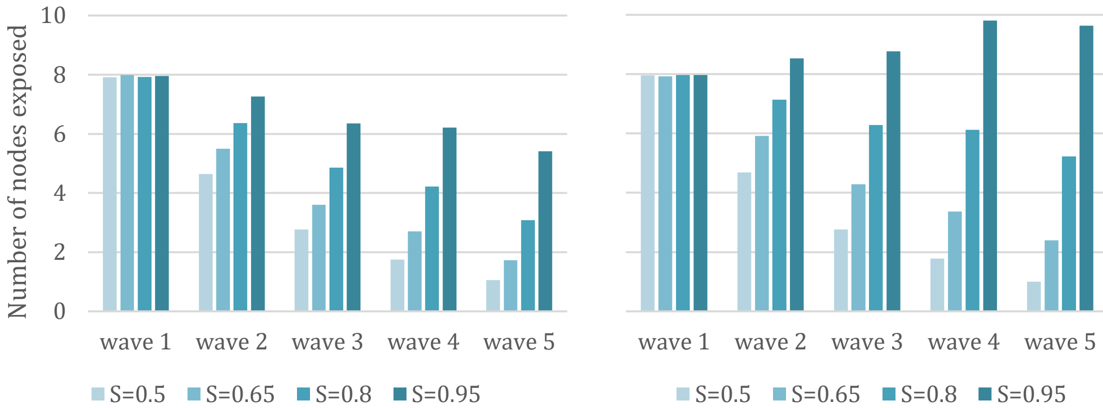

Number of nodes exposed to messages at each wave of diffusion, looking at messages shared at different levels of selective curation (S) in networks of 80% and 95% homophily. In examining information sorting earlier we looked at the proportion of liberal to conservative messages nodes were exposed to. Here we see the total number of messages nodes are exposed to, depending on the selective curation nodes have for that message. If we assume extreme messages are shared with higher rates of selective curation (e.g. 0.95), those messages will proliferate at greater rates than neutral messages with rates of selective curation at 0.5 or 0.65. That viral edge is stronger in more homophilous networks.

There is another take-away from Figure 17. Not only do extreme messages proliferate at higher rates than neutral messages, but - as distinct from information sorting which approaches an equilibrium - the disparity between neutral and extreme posts continues to diverge with each successive wave. In other words, as messages diffuse users are exposed to proportionately more extreme messages.

It may not seem obvious why that would be the case; if all nodes re-share a neutral message at a rate of 5%, while an extreme liberal message is shared a rate of 1% by conservative nodes and 9% by liberals, it might seem that both messages would proliferate at similar levels; 5% is, after all, the average of 1% and 9%.

To see why the results above should not be surprising, let us return to infinite networks which, again, allow for mathematical interpretation. First imagine a perfectly polarized world; that is, a world in which liberal and conservative nodes are in distinct networks – with no cross-connections. Let us assume each node has 100 friends. A moderate message in this network (which is really two disconnected networks) will spread at a rate of $2 * (10 * 0.05)^x$ while an extreme message will spread at a pace of $1 * (10 * 0.01)^x + 1 * (10 * 0.09)^x$. From what we know about exponential growth, the final term for the extreme message, $1 * (10 * 0.09)^x$, on its own will vastly dwarf the moderate message as x increases.

But in our homophilous networks, liberals and conservatives don't exist in distinct networks; there *are* cross-connections between the two ideological groups. To what extent will the spread of extreme messages dominate when there is cross sharing between two groups? We could mathematically express the total number of nodes that are exposed to a message at varying levels of homophily and active curation, but the notation becomes harrowingly complex after just two waves of diffusion. Instead, I model the spread of messages in an infinite network using Python code to see the differential effect of homophily and curation on the spread of extreme versus moderate information.

A simple way to see that effect is to ask when (i.e. under what levels of homophily and message extremity) messages will continue to spread and when they will die out. To do so, we must first observe that in predicting when messages go viral or perish, users' general willingness to re-post information and how many friends they have matter. If either parameter is high enough, then messages will continue to spread regardless of homophily or sharing biases; likewise if they are too low, then a message will always peter out. Those two factors combined – general willingness to re-post and number of friends – are what I will call the "ceiling rate of contagion"; that is, the number of users that would re-post a message if there were no homophily or selective sharing bias. If the ceiling rate of contagion is 2 or above, then a message will continue to proliferate forever regardless of how homophilous or biased the nodes are in what they share. This is because in the worst case scenario (from proliferation's point of view) is if a node has equal numbers of liberal and conservative friends and those friends are 50% likely (times their general willingness to re-post) to reshare the message, then one friend (on average) will forward the message at each wave. If the ceiling rate of conversion is less than 1, on the other hand, the diffusion of a message will always eventually die out. This is because there is no level of homophily and sharing bias that will result in at least one friend reposting a message.

The question then is what happens in ceiling rates of contagion between 1 and 2? In figure 18, I present results of simulations for different ceiling rates of conversion, charting when we should expect a message to continue to spread after 10 waves of diffusion at different levels of homophily and curation – and when it will fade away.

As we see in 18f, when the ceiling rate of contagion is two, as predicted, a message will continue to diffuse regardless of nodes' homophily and curation bias for that message. On the other end, when the base rate is one (18a), a message will continue to spread only if homophily and curation biases are at 100%. In between, we see messages continuing to proliferate the greater the homophily of the network and curation bias for that message.

Figure 18*. When Ideological News Will Spread*

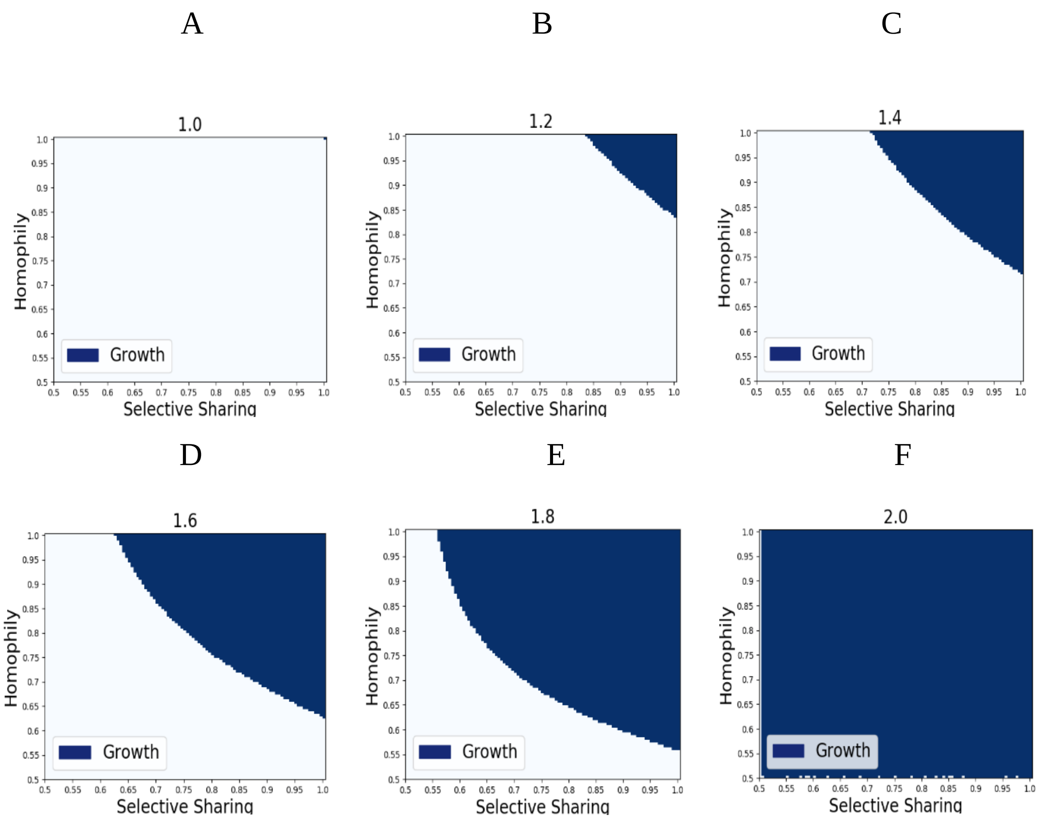

The conditions under which messages will continue to be re-shared in a network ("growth") after 10 waves. Each figure represents simulations with different ceiling rates of contagion (1.0 to 2.0), which is a function of nodes' general willingness to share a message and their number of friends. Messages will always continue to spread at high ceiling rates of contagion (>2) and will always die out if the ceiling rate of contagion is low (<1). In between, the more homophilous the network ("homophily") and the more biased nodes are in sharing a message ("selective sharing") the greater the chances that message will continue to spread.

Again, if we make the assumption that users are more biased in sharing extreme messages, then – all else being equal – extreme messages will have an edge in proliferating through a network, particularly if a network itself is highly homophilous. We see this to be the case in our agent-based model simulations and in infinite models.

From Models to Reality

As the above discussion may have convinced you, the creation of echo chambers and spread of extreme views on social media are complex matters. But as the models in this chapter show, within that complexity some patterns emerge. For one, there is a limit to how much diffusion can sort political information and leave users in ideological bubbles. That limit is the degree of homophily in the network; if 90% of users' friends share their ideology, information will be less than 90% sorted.[26] The extent to which diffusion can promote extreme information, however, is unbounded; if we assume users tend to re-share extreme information, then extreme news stories will drown out moderate ones the more news diffuses.

Outside of those two general observations, diffusion's ability to filter our information bubbles and propagate extreme information depends on the specific parameters of a given network. For diffusion to increase information sorting, users would need to have an ideological bias for *re*-sharing information that is at least as high as their bias in what they initially post. Exactly how much higher depends on users' initial posting bias and on how homophilous the network is. Similarly, for diffusion to amplify the spread of extreme information, it would be necessary that users have a bias for re-sharing extreme information. To gauge how much of an extreming effect to expect, however, we would need to know how strong of an extreme curation bias users have, how homophilous the network is and how much extreme information users are initially exposed to. In short, to know if diffusion amplifies sorting and extreming in an online social network we need to know at least two things about that network: how homophilous is it and how biased its users are in sharing and re-sharing information.

The next two sections are devoted to understanding what those parameters are in real online social networks. In the next section I review existing data on levels of homophily in social

---

[26] Again, this appears to be true given the modelling in the previous section. I, however, offer no proof that this is true for every type of network; I leave open that there exist black swan networks in which sorting can exceed the level homophily in a network.

networks on and off line. Following, I examine a unique data set that allows us, for the first time, to measure users' curation biases.

Chapter 3. Homophily

In the previous section we used models to better understand diffusion's capacity to strengthen information bubbles and amplify extreme information in social networks. But models only give us a picture of what is *possible* given different assumptions about the real world. Ultimately, we want to know what those assumptions should be – in this case, what are social media users' levels of homophily and curation bias?

In this section and the next, I move from theoretical models to examine users' homophily and curation preferences in existing social media sites. How much homophily *is* there in social media? What *are* users' ideological curation biases for what they re-share? And do users have a curation bias for extreme information?

I begin, in this section, with homophily. Although homophily is not required for extreming to occur, it is as close to a necessary precondition for information sorting as we can get. We would want to know at minimum if political homophily does, indeed, exist on social media. But we might also be interested in *how strong* the tendency is for like-minded folks to connect online; figuring out the strength of homophily will lead to better guesses about how much information sorting and extreming to expect on social media. For one, as we saw in the previous section, a network's degree of homophily acts as a ceiling on how ideologically sorted information can be. The exact level of homophily (along with the sharing behavior of the users) also ultimately determines how much sorting and extreming occurs.

Gauging homophily levels might also inform us whether homophily occurs at meaningful levels. Just about every study of online social networks tell us homophily exists, yet what merits that distinction is simply if users prefer to friend more like-minded friends than not – even if that preference is weak. If, for example, 55% of users' connections share their ideology, we would say there is homophily; yet at such a low level that may seem inconsequential. What, though, would be a meaningful level? If users surround themselves with 95% ideological self-replicas, it is easy declare homophily exists at alarming levels. But what if we find users are in 80% homophilous egonets? Or 60%? Deciding if homophily exists in a meaningful way may be a matter of taste.

One meaningful threshold for our purposes would be whether levels of homophily online are any greater than what is seen offline. Theorists propose many reasons why we may expect homophily to be greater - as well as smaller - on social media (briefly discussed below). Remarkably, however, researchers have not offered an empirical comparison - although they have extensively studied homophily in online social networks, primarily Twitter and Facebook.[27]

In this chapter, I review that - at times conflicting – set of research to lay out a range of homophily levels that are known to occur in online social networks. I then compare those levels to our offline lives, likewise using existing research as well as by analyzing data from two national surveys, the General Social Survey (GSS) and the American National Election Studies (ANES). I find a wide range of estimates for our online homophily, one that does not look dissimilar from our offline social networks.

Some Expectations from Theoretical Work

There are a few reasons to suspect our online networks are more homophilous than our real-life networks. We are now familiar with the inevitability of homophily offline; our choices about where to live, work, volunteer etc. will generally surround us with people who share our preferences. We will likewise tend to gravitate towards those similar to us and in turn be influenced by those connections. Yet in real life, we do not have full control over who we encounter day to day. Unless you work at a policy think tank or advocacy group, some of your

[27] Other meaningful thresholds might exist: one would be threshold beyond which negative consequences occur - for example, that affective polarization or intolerance kicks in. Such research, however, does not currently exist and is beyond the scope of this project.

co-workers are likely to not share your political views. Life may also bump you into counter-ideologues at your child's school, your gym or – as the cliché goes - at your Thanksgiving dinner table. On social media, in contrast, it is easy to avoid connecting to acquaintances with attitudinally discordant political views - or to "defriend" them when those challenging views are expressed - and create a homogenous network. It is also possible to further homogenize one's online social world by finding and connecting to like-minded people who you wouldn't be able to meet in real life (Bimber, 1998; Sunstein, 2009).

Yet while social media makes greater levels of homophily possible, there are equally plausible arguments for why online social networks would be less homophilous than our offline world. Offline we are likely to encounter a smaller subset of friends and acquaintances who, precisely because of the frequency with which we spend time with them, will grow to share our worldview. We may also more selectively choose who, among those "close ties," we talk to about politics. Online, in contrast, our day-to-day networks are likely to span a much wider net of "weak ties," including high school friends, distant relatives and that guy you met at the dog run who has unconventional views on immigration (Brundidge, 2010). As opposed to offline, where we may be more selective with whom we talk politics, on social media we may more inadvertently bump into political discussions with people we'd likely never have a conversation with in real life (Wojcieszak & Mutz, 2009).

Levels of Homophily Online

As mentioned earlier, there is no shortage of evidence that homophily exists on social media. Without exception, researchers find ideological and partisan homophily on social media sites, whether looking at how much co-ideologues overlap in who they follow (Boutyline & Willer, 2017), who users retweet (Barbera et al, 2015; Conover et al, 2011) or mention (Conover et al, 2011), or how many of users' connections share their ideology (Bakshy et a;. 2015; Colleoni et al, 2014) or share their views on issues (Goel et al, 2010).

The exact levels of homophily online, however, are not as easily gauged looking at existing research. The first obstacle to estimating levels of homophily across social media sites is that researchers do not use similar measures. One of the more frequently cited articles, for example, measures homophily as the increased likelihood that users will mention or retweet a co-ideologue (Conover et al, 2011). Another measures homophily as the likelihood that users' connections will follow the same political actors (Boutyline & Willer, 2017).

Fortunately, there is one measure – the proportion of users' connections that share their political leanings – that is used frequently enough to give us a ballpark estimate for homophily on Facebook and Twitter.[28] Unfortunately, at first glance, it is a particularly large ballpark - 67% to 96%.

What should we make of this broad range? For one, we should expect variation in homophily across platforms, because each are designed for different types of communication and community. Facebook, for example, is designed primarily for real-life connections and requires users to mutually accept friendships. We might expect Facebook to be more homophilous than Twitter which does not require reciprocity and so allows users to connect to a more eclectic array of users and users they do not personally know (Colleoni et al, 2014). We might also expect greater homophily on platforms that focus on political discussion, such as many subReddits, than on non-political topics, such as mommy groups (Wojcieszak & Mutz, 2009).

But inter-platform differences do not explain all the variation in estimates we see. Differences also comes from two choices researchers make in constructing their estimates; first, how to

[28] I restrict my gaze to those platforms for two reasons. One is purely practical: that is where the research is. Most scholars examining political polarization online look at either Facebook or Twitter. They have good reason to do so, since it is on those platforms where political information has its widest reach and largest volume. According to surveys from Pew in 2017 (Shearer & Gottlieb, 2017), 44% of Americans get some news via Facebook, and 12% via Twitter users (as for the other platforms: 4% get news via Reddit, 6% from Snapchat, and 8% via Instagram). Although Twitter's reach is much smaller, it arguably has an outsized influence on the news agenda of more traditional media (Swasy, 2016).

measure the political leaning of users and, often dependent on that first choice, which users to include in estimates.

By and large, social media users do not publicly declare their ideology or partisanship on their social media profiles. Even when given a specific feature to do so, as on Facebook, only one in ten take the opportunity (Bakshy et al, 2015). Without such self-reports measures of ideology come in one of two types; 1. evaluations of the content of users' messages; or 2. evaluations of the ideology of accounts users follow.

Barbera et al (2015) use the latter approach in what may be the most widely cited and adopted method for estimating the ideology of Twitter users (Barbera, 2014). Using a Bayesian spatial model, they infer the ideology of Twitter users based on the accounts of political actors that they follow. Using this measure, the authors finds that 67% of the median user's connections are co-ideologues.[29]

Colleoni et al (2014) likewise measure users' ideology on Twitter, but do so by analyzing the content of their tweets. Their estimates for homophily not only differ from Barbera et al (2015), but also vary widely depending on the population of users they analyze. Drawing on the full Twittersphere in 2009, the authors first measured the homophily of users who had shared at least one political tweet[30] and found that users were, on average, 80% homophilic. This number, however, hides a few complicating dimensions. Because the network of Twitter users they identified were overwhelmingly Democratic, if users randomly connected to other users, average homophily would be 77% - which makes it seem users have hardly any bias toward homophily at all. Yet when the authors separate out partisans, they find that Democrats would connect to co-partisans 79% of the time at random, but in reality do so 88% of the time. Republicans, in contrast, would connect to other Republicans 63% of the time if at random, but do so only 23%

[29] Barbera does not distinguish varying levels of homophily between liberal and conservative Twitter users.

[30] The researchers used a machine learning method, basing their training set on tweets of users who exclusively followed either Democratic or Republican politicians on Twitter.

of the time.[31] In other words, Democrats appear to have a bias toward homophily while Republicans do not. In part because the authors were surprised by the anti-homophilic behavior of Republicans they took a second pass at measuring homophily, this time only looking at users who follow political figures on Twitter. Limiting themselves to this smaller pool, their results flip; random connections would predict Democrats to be 51% homophilic and Republicans 93%, but they are in fact 44% and 96% homophilic respectively. These results should make us particularly cautious when considering measures of homophily; depending on which subset of users we include in the homophily pot, we could get vastly different results.

Because of its strength in identifying user's political leanings, I take the Barbera et al (2015) estimate as a more accurate representation of the homophily of Twitter users.[32] But should that estimate be an upper or lower bound? Since their estimate is limited to users who follow political accounts, and since we know that the politically engaged tend to have stronger views (Mutz, 2002b), we might expect that the Twitter users in their set would be more selective in connecting to others. Other analyses in their paper suggest this is the case; users they identify as "strong" ideologues have more homophilous egonets. It thus is likely the users in their estimates - who all follow at least one political account - are more homophilic than the average Twitter user.

Turning to look at homophily levels in Facebook, the most cited estimate comes from Facebook's own research team who have the advantage of having access to the self-identified ideology of 10.1 million of their active users. The authors report that among those users, on average 18% of liberals' friends were conservative and, similarly, 20% of conservatives' friends were liberal; in other words, liberals appear to be 82% homophilic and conservatives 80% homophilic on average. These proportions seem more homophilous than the findings of Barbera

[31] The expectation of 63% homophily - when Republicans only comprise 10% of the population - may seem surprising. If the researchers were using purely random attachment we would expect 10% homophily. Unfortunately the authors don't say how their random graph was constructed to reach this number.

[32] Barbera (2015) validates those estimates by cross referencing to campaign donation records (>90% accuracy) and voter registration records (>75% accuracy).

et al (2015), but they are not directly comparable primarily because moderates are included in Facebook's denominator, whereas Barbera et al (2015) only considered connections among liberals and conservatives. If we were to convert Facebook's estimate by removing moderates from the mix, homophily drops to 77%.

The Twitter and Facebook estimates of homophily we have looked at so far are, again, based on users we can assume are more politically active - as they either follow a political account, have tweeted a political message or have bothered to identify their political leaning. Again, this is a special subset of users that we might expect to have more homogeneous ties. To get a full picture of levels of homophily on social media we would need to study the networks of the politically inactive as well. The challenge, as we have seen, is identifying the political leanings of social media users who don't engage in politics (by posting, linking or self-labeling).

Goel et al (2010) partly get around this limitation by launching a Facebook app that asked users to take a survey about their political beliefs - and then using those answers to measure homophily. Although users who choose to download the app may be a self-selecting group who, again, is likely to be more politically engaged, the app was advertised through multiple avenues and so had the potential to reach users who might not be normally politically engaged online. Of the 2,504 users surveyed, the researchers identified over 12,000 friend pairs between 900 users. On average, they found, friends agreed 75% of the time (across issue topics and ideology). But adjusting for the fact that a random pair in this sample would agree 63% of the time, the authors surmise that for an issue where there would be overall 50/50 agreement, friends would be likely to agree 67% of the time. While this number looks similar to the Twitter estimates of Barbera et al (2015), there is a key difference; Goel looks not at ideology or partisanship but at agreement across policy issues. Since Americans tend to be less polarized on issues (Fiorina et al, 2005), we may not be surprised that homophily on issues is weaker than ideological homophily on the same platform.

It should also be noted that the measures above are imperfectly comparable for one more reason: some calculate homophily at the individual level and others at the population level. Goel et al (2010), for example, does the latter, simply looking at the proportion of all 12,000 connections that are congruent or incongruent. Barbera et al (2015), in contrast, looks at each individual's level of homophily and then averages across those individuals. Assuming not all nodes are connected to each other, calculating at the individual or population level will result in slightly different estimates of homophily. To see how that might be the case, consider Figure 19a where there are a total of four connections, two congruent (solid line) and two incongruent (dashed line). At the population level, homophily is 50%. At the individual level, however, we see there is one node with 100% homophily, one with 2/3rds, one with 1/2 and one with 0, resulting in an average homophily of 54%. Unfortunately, without full access to the researchers' data sets, it is impossible to recalculate the measures above or even know if, say, we should expect homophily at the individual level to be higher than homophily at the population level. That would be the case if nodes with fewer friends tended to be more homophilic; yet while plausible we do not know if that is the case. The best we can say is that the two measures will not be far off from each other; to construct a network where the two measures diverge significantly (for example in Figure 19b where population homophily is 50% and average homophily is 58%) one has to be creative in structuring a network. We might then be confident that both measures fall within a couple of percentage points of each other.

***Figure 19**. Two Ways to Calculate Network Homophily*

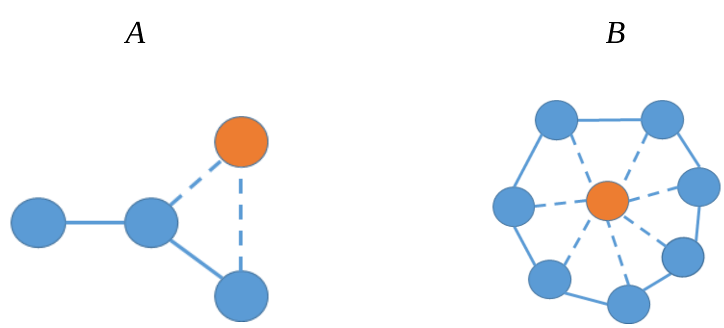

These micro networks demonstrate the differences between calculating average homophily and network-wide homophily. In A, average homophily is 54% while network-wide homophily is 50%. In B, 58% and 50%, respectively.

Such small discrepancies, moreover, unlikely account for the wide range of homophily estimates we see. Depending upon how we categorize the leanings of users, which users we include and the nature of the platform, we see levels of homophily anywhere from 67% to 96%. Yet, if we set aside Colleoni's estimates as outliers, a narrower range - and more cohesive picture emerges. At one end, on Facebook, approximately 77% of users' friends share ideologies while on Twitter the average is 67%. The higher estimate on Facebook might be expected because, again, these are mutual friends and so are more apt to be close ties. Both estimates, however, may be upper bounds, as these samples are likely to be more politically engaged than the average social media user. Goel's study, which is not limited to the politically engaged, suggests that Facebooks' users may have yet more diverse friendship networks - with homophily as low as 67%. If we assume that expanding the breadth of users to the less politically active would diminish estimates of homophily, then Twitter users may themselves be less than 67% homophilic on average. Given an upper bound of 77% homophily and a lower bound less than 67%, a conservative estimate of homophily on social media might be in the range of 70-75%.

Homophily in Real Life

How does a 70-75% homophily estimate on social media compare to homophily in real life? While estimating homophily online presents challenges, they pale in comparison to challenges of estimation offline. On social media the primary hurdle is identifying the ideology of users. Once that is known, however, it is then trivial to observe who connects to whom. In assessing whom we associate with in real life, by contrast, researchers generally have an easy time identifying the ideology of individuals (by asking them in surveys), yet have difficulty observing links between those individuals.

The standard study on real life homophily depends on self-report, asking individuals to identify their ideology and the ideology of those they interact with at various levels of intimacy.[33] This approach has its weaknesses: subjects may not only have difficulty recalling who they associate

[33] Not all studies. In "snowball" studies run by Huckfeldt, about to be discussed, researchers survey subjects' friends as well in order not to depend on their friends' estimates. Yet other studies, usually conducted in organizations, are able to create a fuller map of connections by interviewing all members or by using technology to observe who connects to whom (e.g. Orbach et al, 2015).

with, but may also have a poor sense of what their friends' political leanings are (Huckfeldt & Sprague, 1987). If survey participants were simply inaccurate at recalling or guessing their friends' ideologies, this at least would not be an inferential problem for researchers; subjects' errors would cancel each other out giving us a good approximation of the average level of homophily. The risk, however, is that subjects' errors are not random, but instead share the same biases or cognitive shortcomings. One such known bias is that subjects will overestimate how much their friends share their political views (Huckfeldt & Sprague, 1987). Another - conflicting - bias could be that subjects will want to believe their social networks are more diverse than they are; nobody likes to admit they're a closed-minded ideologue (Klar & Krupnikov, 2016).

As with estimates for homophily online, measurements for offline homophily also come in different flavors. Yet a couple of well-regarded papers use a measure homophily offline similar to studies above by simply looking at the proportion of like-minded connections.

Gentzkow and Shapiro (2011) use data from GSS to estimate offline homophily as a proportion of connections that share one's ideology.[34] The GSS survey, given in 2006, asks subjects to think of people they are "acquainted with" in their neighborhood, at work, and in voluntary organizations[35] as well as to think of people in their family and people whom they trust. Subjects are then asked how many they are "pretty certain" are "strongly" liberal and conservative, given the choices 0, 1, 2-5, 6-10 or more than 10. Looking at the answers of all subjects who identify themselves as at least "slightly" liberal or conservative, and converting those bin answers to 0, 1, 3.5, 8 and 12, we can see in Table 1 that homophily ranges from 58% (for acquaintances at voluntary organizations) to 66% (for people subjects say they trust):

[34] This is not, however, the ultimate measure they use in their paper. Rather they measure "ideological segregation" as the distance between the proportion of connections who are conservative among conservatives and liberals.

[35] Such as "schools, clubs, associations, or places of worship."

Table 1*. Homophily in Offline Networks*

	All ideologues	**Solid ideologues**
Voluntary associations	57.8%	60.6%
Neighborhood	60.2%	63.7%
Family	62.6%	66.3%
Acquaintances	62.5%	68.7%
Work	58.5%	60.8%
Trusted friends	65.5%	70.4%

Levels of homophily (% of connections that are strong co-ideologues) among social groups for all ideologues and strong ideologues. GSS 2006.

Those levels of homophily seem substantially lower than what we saw on social media, but there is at least one reason we would expect them to be. Recall that the estimates for homophily on Twitter and Facebook were restricted to presumably more politically engaged users and thus users with stronger ideologies. The GSS sample, in contrast, includes people with weak as well as strong ideological leanings, so we would expect them to have more diverse networks. If, however, we restrict the GSS sample to only those with relatively strong ideological identification (by removing subjects who noted they were "slightly" liberal or conservative), levels of homophily range from 61% (voluntary associations) to 70% (trusted friends).

Even with this adjustment, we should be cautious in using these estimates for reasons already mentioned. Subjects are known to have less than perfect recall when asked to describe their worlds. People are subject to an array of limitations and biases when asked to recall information; they may pull on information that is most accessible - because it is recent or salient - that then gives an inaccurate representation of their lives (Schwarz, 1999). Most critical to our concerns, subjects may want to answer in ways that project a positive image of themselves. If, for example, subjects wanted to think of themselves as open minded (and not stuck in an echo chamber), they may imagine they associate with more counter-ideologues than they do. Their answers, however, could also suffer from the opposite bias; they may overestimate how much others agree with them and thus underestimate how many counter-ideologues are in their circles (Goel et al, 2010; Huckfeldt & Sprague, 1987).

Huckfeldt et al (1995) uses a survey methodology that helps mitigate those potential biases by forcing subjects to be more specific and concrete in thinking about the ideology of their friends. Looking at levels of agreement among citizens in Germany, Japan and the US in 1992, researchers asked participants for up to four names of people they talk to about important matters and then one more name of someone (not on that list) that they talk to about politics. Participants then note the party those acquaintances prefer. Measuring the proportion of all connections that agree at the population level (as in the Goel paper), they find that among the acquaintances whose party preferences US participants can identify there is 65% homophily.

The 2009 ANES survey uses a similar approach, asking participants to list the first names of up to eight people with whom they spoke “about government or elections” in the previous six months. They are then asked (for up to the first three names) whether that person “probably think[s] of him or herself as a Republican or Democrat.” Among 1524 self-identified partisans who said they had spoken to at least one person about politics, on average 79% of their partisan interlocutors shared their partisan identity. If we include in this analysis subjects who “lean” Democrat or Republican, the rate of homophily decreases to 72%.

This estimate is higher than the ones we see in the GSS data and Huckfeldt et al (1995) study, but the reader likely has already surmised why that would be the case. In the ANES survey subjects are explicitly asked to think about whom they discuss politics with. It is reasonable to guess the talking partners who come to mind are those the subject speaks to most frequently about politics and thus are more likely to be close connections that share their political beliefs. The GSS survey, in contrast, first asks subjects to think of a wide array of acquaintances, which would include acquaintances subjects are unlikely to discuss politics with. The prompt in Huckfeldt et al (1995) about “people you trust” similarly might capture a more ideologically diverse crowd than the ANES’s “people you discuss politics with.” Another explanations for the divergence could be the difference between asking about ideology (GSS and Huckfeldt et al) or partisanship (ANES). Finally given trends affective partisanship, it is also possible that part of the difference between the Huckfeldt et al (1995) and ANES estimates is due to rising homophily (among discussion partners) between 1992 and 2009.

It must also be noted that the ANES question does not necessarily exclude discussants on social media. Although the survey taker asked subjects if they “talked” to anyone, they left open the medium of discussion to include “face-to-face, on the phone, by email, or in any other way.” It is more than possible, then, that respondents’ lists include friends that they “talk” to on social media.

Homophily On and Off Social Media: The Upshot

So are our networks on social media more homophilous than our offline networks? The evidence above is inconclusive. Reliable estimates for homophily on Twitter and Facebook range from 67-77%, but those estimates are based on networks of more politically engaged users, so we can assume they are higher than for the typical user. Offline, if we restrict estimates to more engaged individuals, we see estimates ranging between 61-70%. If, however, we focus on a small subset of individuals’ most frequent political interlocutors, homophily rises to 79%.

The only thing the evidence may tell us is that we have no strong evidence that our friends on social media are any more homophilic than our friends in real life. Likewise, we would be hard pressed to say that our offline friendship networks create more of a bubble than our social media networks. If anything, the estimates above suggest that homophily online is a near reflection of our offline worlds. This would align with theorists who argue that technology - in particular innovation in communication technology - does not change human behavior, but rather gives humans another platform to extend that behavior (Gaver, 1991; Gibson, 1977).

Moving forward in this dissertation, however, we can take away a good working estimate for homophily online of about 70%. This is at the lower end of the 67% - 77% range yet assuming that range, again, overestimates average homophily on social media a 70% estimate is possibly conservative. In the final section of Part I, I will use that estimate to re-run the models in the previous section. But first it is necessary to get an estimate for users' curation biases on social media. I do so in the next chapter, using data drawn from Twitter.

Chapter 4. Curation

In 1949, David Manning White, a journalism professor at a Midwest university, asked a local "wire editor" to do something that had never been done before - hold on to the wire stories he decided not to print (White, 1950).

White was exploring an idea that to become news, events had to pass through a series of gates manned by "gate keepers" each deciding which stories merit being published. A newspaper's wire editor acted as a final gatekeeper; as stories came through the three major wire services - at the time, Associated Press, United Press and International News Services - the wire editor chooses which to run in the next day's publication and which to put in the trash bin. At White's request, instead of discarding stories that didn't make the cut, for one week this wire editor agreed to hold them in a box for the young professor.

White knew that to understand why some new stories made it through the gate, you need to see all the articles that vied for publication. Anyone can see the stories that make it into print; White could now see the full set of stories the editor chose from – and so see which factors he selected for.

Similarly today if we want to measure political biases that affect social media users' choices about what to share, it's not enough to see what they "print" on their walls – we also need to see the stories they pass over.

Existing research on social media users' curation behavior only looks at their "printed" stories and leave out the set of stories users choose from. Those studies confirm the intuition that ideologues and partisans tend to post news stories that align with their beliefs; that is, they are *de facto* ideological curators (An et al, 2014a; An et al, 2014b; Barbera et al, 2015; Shin & Thorsen, 2017). But we do not know if users also *actively* select politically congruent news to share with their friends, or if they simply pass on any given story on their feed with some constant probability. We also do not know if they have a preference to re-post extreme information.

Therefore, as with understanding what becomes news, we need data that shows us not only what news users post, but also what stories they choose from. I collect such a data set from the social network site Twitter.[36] I choose Twitter for the same reason most researchers do: unlike Facebook and other online social networks, Twitter's data is publicly available. Anyone with a Twitter account can use the platform's API to query and download its data, although with limits.[37]

Like all online social networks, Twitter has its own unique set of characteristics which shape the culture, norms and behavior of users; for example, text used in posts is limited to 280 characters and connections between users do not have to be mutual. Twitter users are also more educated and more male than users of other popular social network sites including Facebook, Instagram and Snapchat (Shearer & Gottfried, 2017). As such, we cannot assume that the sharing of behavior of Twitter users represents the behavior of all social media users.

At the same time, Twitter may be a particularly apt platform to study if we are interested in how political information is shared. Twitter users tend to be more politically engaged than average (Shearer & Gottfried, 2017) so while Twitter as a whole may not reflect what occurs on other platforms, it may represent the exchange of political information that occurs on those platforms. Twitter also plays an outsized role in the dissemination of political news, often influencing the news picked up by news organizations; so we may be interested in Twitter as a phenomenon in its own right.[38] Finally, it may be that Twitter's specific features, its "affordances" in the

---

[36] The data I collect allows me only to look at users' selective curation from among the tweets that they see; I do not collect the totality of news that users are exposed to from all media. Yet, as will be discussed, even limited to observing users' choices from among tweets we can make educated guesses about their selective curation from other media sources.

[37] The freely available API is known as the "garden hose" because of the relatively restricted flow of data one can pull from. There is also a paid-for "fire hose" service for collecting data which this researcher did not use.

[38] 74% of its users receive at least some news on the platform (Shearer & Gottfried, 2017) and virtually all media companies and political actors have a presence on Twitter.

language of media studies, and its norms are what make it fertile ground for news dissemination. Thus we might suspect that future social media platforms that are popular for sharing political news will have characteristics in common with Twitter.

Data & Measures

To measure users' active ideological curation, I collected the Twitter posts of a convenience sample of 472 Twitter users over a three month period from September through November, 2016. During that period I also collected the tweets from a sample of the Twitter accounts those users followed.

Most of those 472 Twitter users were recruited from Amazon's Mechanical Turk. As participants in previous online surveys they were asked if they had a Twitter account and, if so, if they would share their Twitter handle and political ideology, and also allow researchers to observe their tweeting behavior. Among those who said yes, I followed all who had active accounts - i.e. who had posted at least one tweet in the past six months.

Just as Twitter users may not be typical social media users, Mechanical Turk workers also have atypical demographics. Mechanical Turk workers are more male (54% as opposed to 48%), younger (32, 48), more educated (45% with a college degree instead of 29%) and more liberal (46% identify as Democrats as opposed to 36%) than a representative sample of American adults. Yet some characteristics that distinguish Mechanical Turk workers from the US population make them similar to Twitter users who are likewise more likely to be male (53%), younger (42) and highly educated (45% with college degrees). The exception is partisanship; 36% of Twitter users as compared to 46% Mturk workers identify as Democrats. Perhaps the greatest distinguishing factor between Mechanical Turk workers and Twitter users is race; 59%

of Twitter users identify as White while 72% of Mechanical Turk workers do – and the latter number is possibly an underestimate (Levay et al, 2016; Shearer & Gottfried, 2017).[39]

In addition to the 472 panelists, I also collected the outgoing tweets and a sample of incoming tweets for 5000 random Twitter users.[40] Unlike the 472 panelists, I did not have the self-reported ideology for these users, but their data can be used for examining extreme curation. While it is necessary to know the ideology of users to measure their ideological curation biases, no such requirement is needed to measure their bias toward sharing extreme information; one merely needs to know the ideology of tweets users are exposed to and that they re-tweet.

User Ideology

For user ideology I use self-reports of the 472 panelists, each of whom placed themselves on a seven-point scale from very liberal (1) to very conservative (7). This measure improves upon previous work on curation biases on Twitter, which usually infers user partisanship and ideology from indicators on their Twitter accounts, such as the content of users' tweets or the political accounts they follow. While inferred measures of political leanings may be reliable (of note, Barbera (2015) has greater than 90% validity), they confound our ability to test for both de facto and active curation. Estimates of user ideology that depend on the accounts users follow, for example, may be correlated with the ideology of users' tweets for two reasons: they may, indeed, correlate because a user has an ideological curation bias, but it is also possible because the tweets a user selects its re-tweets from will be from the accounts they follow. If users' ideologies are inferred from those accounts, we should not be surprised that a user identified as conservative also posts conservative tweets. Even more problematic are measures of user ideology that depend on the content of the user's tweets, in which case the measure for message content and the user's

[39] In the Mturk study conducted by Levay et al (2016) 7% of respondents did not identify their race.

[40] To collect a random sample of Twitter users, I used a random number generator (in Python) and used Twitter's API to check if there was a user ID associated with that number. Among existing accounts found, I collected tweets of 5,000 of users whose accounts were English language and who had tweeted at least once in the previous 6 months.

ideological leanings will be nearly identical. Self-reported ideology, in contrast, gives us an independent measure that lets us avoid the use of collinear and possibly confounding variables.

Tweets and Tweet Ideology

To estimate the ideology of the political information that users post and are exposed to it is necessary to address two challenges. For one, there must be a way to determine if a tweet is "political." Second, we must reliably and validly code political content as liberal or conservative. To kill two birds with one methodological stone, I use the ideological news index created by Bakshy et al (2015) which estimates the ideology of 500 news sites based on the sharing behavior of 10 million Facebook users with self-identified political leanings. To cull out soft news sites, the measure uses a trained "support vector machine" to identify only hard news. I use that index to assign an ideology to all tweets that include a link to one of those 500 sites. A shortcoming of such a measure is that it is based on the behavior of users a year before the tweets in this study were collected, during which time there may have been shifts in the ideology of news sites. Assuming any shifts were randomly distributed (i.e., there was no imbalanced trend), however, any changes from the time the Bakshy index was created should not bias these results.

I likewise use the Bakshy measure to assess the ideology of tweets users are exposed to. Ideally we would measure every tweet from all accounts followed by those 5,472 users, but given the limits of data collection set by Twitter's API I instead selected a random sample of 100 of the twitter accounts each user followed, collecting all the tweets from those accounts over the same three month period.[41]

We must be cautious, however, in saying that this represents the tweets users are exposed to. In February 2016 Twitter began promoting the posts it believed its users, based on their previous behavior, would be more interested in (@mjahr, 2016). This type of algorithmic filtering is

[41] Tweets were collected at the end of each month using Twitter's API and the Python "Twitter" package. Because Twitter limits the collection to the 3,200 most recent tweets it was necessary to collect the tweets of prolific accounts twice a month.

known to result in greater exposure to ideologically consonant information in Facebook (Bakshy et al, 2015) so we might assume it would have a similar effect in Twitter. Given that is the case, the estimates we find for users' exposure to ideologically friendly information are likely underestimates.

In using the Bakshy measure, I am both able to limit my analysis to the spread of hard news stories while also having a reliable estimate for the ideology of those news stories. The downside is that by using only links to hard news sites, I limit the number of tweets in the data set. Even so, I end up with workable number of users and tweets; over the three-month period 140 of the panelists and 571 of the random users posted at least one link to a news site in the Bakshy index, posting a total of more than 6,000 links.

Analyzing Ideological Curation

To determine users' ideological curating biases, I look at the tweeting behavior of the 472 panelists only, as these are the users for whom we have self-reported ideology. Given what we know about Mechanical Turk workers it is not surprising that our sample leans liberal (Figure 20). Among the 140 who tweeted out at least one political link, 89 report being liberal, 30 as moderate, and only 21 as conservative. Because of the relatively small number of conservatives in this sample, I mostly focus on liberal users.

Figure 20. *Ideology of Twitter Users in Sample*

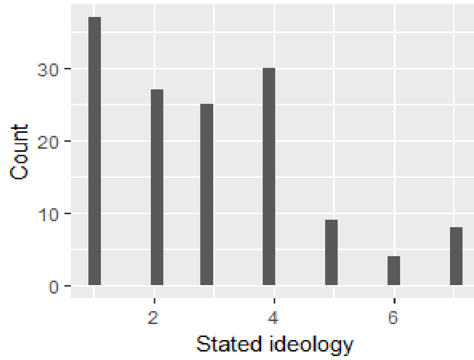

Count of subjects in panel by stated ideology from very liberal (1) to very conservative (7).

To assess levels of *de facto* ideological curation (again, how much ideologues tend to share ideologically consonant information, regardless of what tweets they themselves are exposed to), I use two measurements: the *average ideology of the tweets* users post and the *proportion of liberal to conservative tweets* shared by users, both based on the Bakshy estimates.[42]

As we can see from Figures 21a and 21b, users' stated ideology is, as expected, associated with the ideology of tweets users post both when looking at average ideology ($p < 0.01$) and, to a lesser degree, at the proportion of liberal tweets ($p < 0.1$). The average tweet ideology of strong liberal users (ideology = 6 or 7) is 0.33 on a scale from most conservative (-1) to most liberal (1), while 80% of their tweets on average are liberal leaning. Moderate users (ideology = 3, 4, or 5), by comparison, have an average tweet ideology of 0.19 and a 67% liberal tweet rate. Strong conservative users (ideology = 6 or 7), finally, have a 0.06 average tweet ideology and a 63% liberal tweet rate.[43]

---

[42] I reverse code these estimates so that it is easier for the reader to visually compare "average tweet ideology" with "proportion of liberal tweets" in the figures; otherwise the trends would appear at first glance to go in opposite directions.

[43] Recall that there are only 12 conservative users in this sample so their tendency to post liberal tweets may be an aberration.

Figure 21. *De facto Selective Curation: The Association Between User Ideology and Ideology of their Tweets on Twitter*

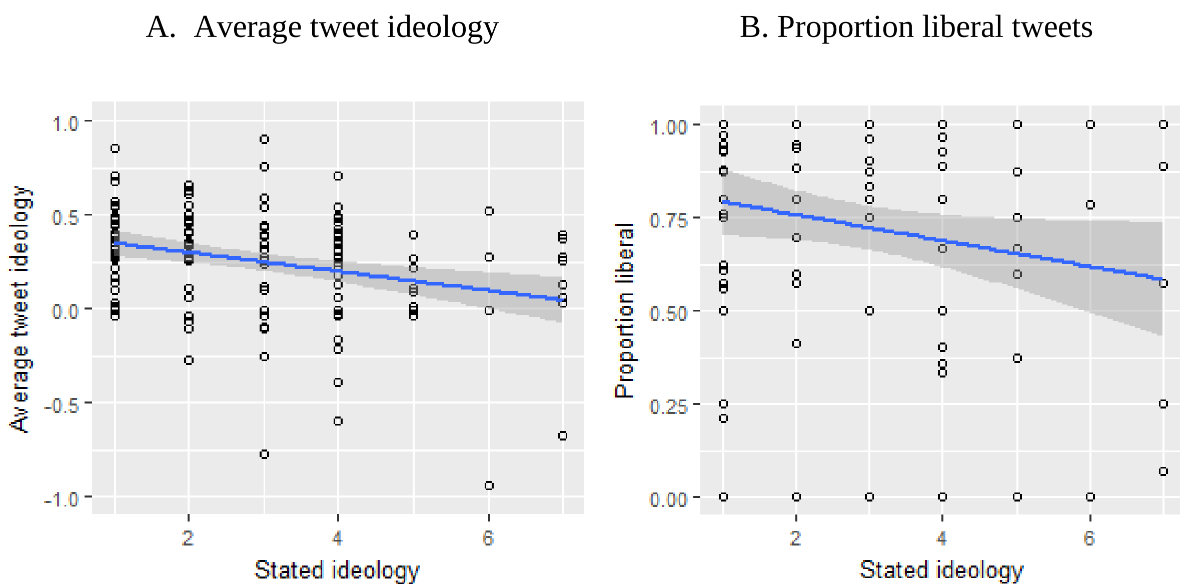

X-axis: stated ideology of users from (1) very liberal to (7) very conservative. Y-axis: (a) Average ideology of posted tweets using reverse coded Bakshy Index from very liberal (1) to very conservative (-1) and (b) proportion of liberal tweets (positive on Bakshy Index). OLS regression lines w/ 95% CI shown.

Even though this sample confirms previous findings – and common intuition – that the ideology of a social media user is associated with the ideology of what they share, it should be noted that there is considerable cross-ideological sharing. Even those who call themselves "very liberal" are not purists; on average 18% of the news links they post are to conservative sites.[44]

What of *active* curation? Recall for us to believe active ideological curation occurs, we must observe that tweets users post are more ideologically-homogeneous than those they are exposed

[44] The reader may wonder if at least part of this "cross-ideological" sharing is a product of the "retweeting is not endorsement" phenomenon; that is, a retweet of a NYTimes article could mean one shares the sentiment of that article, or – in the feed of a conservative – could come with the presumption of mockery. While it is impossible to know the sentiment of a straight retweet (that is, a retweet with no annotation from the retweeter), one can look at "original" tweets that link to articles and include a comment from the poster to see if the user's comment indicates mockery (or another form of non-endorsement) of the article it is sharing. To see how pervasive non-endorsement tweets might be, I hand coded 500 such tweets and was only able to identify 13 (2.6%) as likely non-endorsements. Although a data analysis with a small sample and one coder, it indicates that non-endorsement tweets are unlikely to account for all the cross-ideological posting observed among Twitter users.

to. As a first step toward detecting active curation, I compare the average ideology of tweets users are exposed to (“incoming”) alongside those they post (“outgoing”), and likewise compare the proportion of users’ incoming and outgoing liberal tweets. As seen in Figure 22, there is little difference between the tweets that come in and those that go out; if anything users appear to have an inverse bias, with liberals posting links that are less liberal than the tweets they are exposed to.[45]

Figure 22*. Ideology of Incoming and Outgoing Tweets*

A. Average tweet ideology B. Proportion liberal tweets

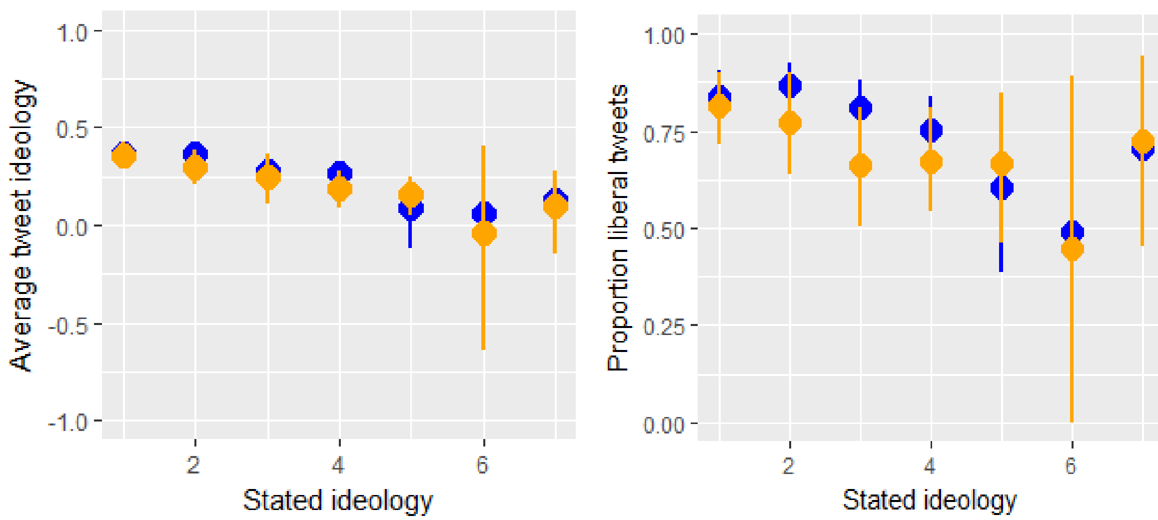

Average ideology of tweets (A) and proportion of liberal tweets (B) that users are exposed to (blue) and that they tweet out (yellow), looking across very liberal (1) to very conservative (7) Twitter users in panel sample.

The graphs in Figure 22, however, compare incoming tweets to *all* outgoing tweets; it may be more apt to look just at *re*tweets and disregard “original” tweets. Retweets are, as they sound, tweets that a user re-posts on their wall, where they appear along with a notation that it is a retweet (“RT”) and a reference to the originating user. Users also have a choice to add a

[45] We should also recall that Twitter likely prioritizes ideologically congruent tweets, so the estimate for incoming tweets may underestimate the proportion of ideologically congruent tweets a user actually sees.

comment to the post; in Twitter jargon, a retweet with a comment is an “annotated” tweet. Original tweets are generated by the user either by copying and pasting an outside link into a twitter post or by using a news site’s Twitter app. Of the two types of tweets, retweets give us a more precise picture of bias in curation, since we can more directly compare the set of information users are exposed to and what they choose to retweet. Original tweets, by contrast, would need to be compared to the totality of the news environment users are exposed to both on and off Twitter, information this researcher does not have access to.

Zeroing in on retweets, Figure 23 shows the ideological difference between what users see on their Twitter feeds and what they retweet. Again, there is little movement; users, on average, retweet what they are exposed to without seemingly discriminating between left and right. The possible exception are those who identify as just left of center; their outgoing retweets are slightly more liberal than those in their Twitter feed.[46]

---

[46] The propensity of liberal centrists to retweet proportionately more liberal links is seen more strongly when just looking at tweets from “unverified” accounts. That centrists are the only ones to exhibit a curation bias may, at first, seem surprising since we might expect more extreme users to have a stronger bias in what they re-share. A plausible explanation, however, is that we are seeing a ceiling effect: extremely liberal users can’t retweet more liberally than what they are exposed to because their Twitter feeds are so liberal to start with.

Figure 23. *Ideological Difference Between What Users See and Retweet*

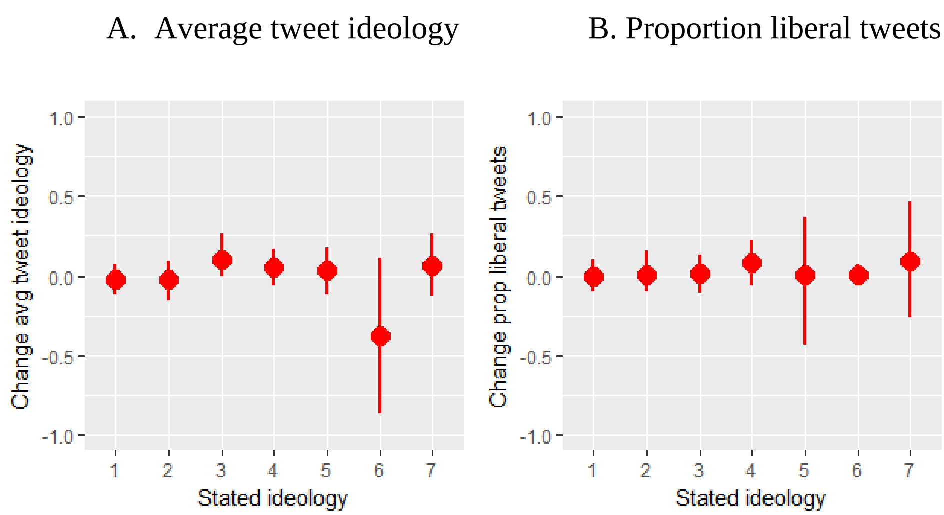

Difference in average ideology of tweets (A) and proportion of liberal tweets (B) between what users are exposed to and what they retweet.

Figures 22 and 23 give us visual evidence that users are not biased curators. To confirm *statistically* that there is no difference between users' incoming tweets and outgoing retweets, I run a multiple linear regression of user ideology on outgoing tweet ideology holding constant the ideology of a user's incoming tweets. Whereas we saw previously that user ideology has a strong association with outgoing tweet ideology on its own, when accounting for ideology of incoming tweets that relationship is weakened when the unit of analysis is average ideology ($p = 0.08$) and becomes statistically insignificant when comparing the proportion of ideologically-congruent tweets ($p = 0.31$). As Figure 23 suggests, the relationship is even more tenuous when looking only at retweets ($p > 0.1$ for both average ideology and proportion of liberal tweets).

One Explanation for the Absence of Active Ideological Curation

This absence of biased retweeting contradicts popular intuition. A possible explanation why we do not see biased curation, however, is that social media users are adept at *pre*-selecting their information environment. If users choose to follow news sources that align with their ideology, they obviate the need to impose a selection bias on the stories they share. In other words, it may be that any ideological selection bias happens at the level of selective exposure (users' bias in what accounts to follow) and not selective curation.

If this were the case, the ideology of incoming tweets should mediate any effect of user ideology on their tweets. To test for such a relationship, I use the R mediation package of Tingley et al (2014). Results, shown in Figure 24a, demonstrate that, indeed, any effect of user ideology on retweet ideology is mediated through the ideology of incoming tweets (Average Causal Mediation Effect or ACME), whereas ideology has no direct effect (Average Direct Effect or ADE).

That mediated relationship becomes even more pronounced if we distinguish between "verified" and "unverified" accounts that users follow. Verified accounts are those that Twitter deems to be of "public interest" and that are "typically maintained by users in music, acting, fashion, government, politics, religion, journalism, media, sports, business, and other key interest areas" (Twitter). We can assume that verified accounts are more likely to belong to figures whom users follow due to their public presence rather than due to a social connection. Unverified users, in contrast, are more likely to be followed because users are friends or have some other kind of offline relationship (Barbera, 2014).

If the reason we do not see biases in retweeting is because users are, indeed, carefully selecting their information environment to reflect their ideology, we would expect a stronger mediating relationship when looking at retweets from verified accounts than from unverified accounts.[47] Figures 24b and 24c show evidence to support exactly that conjecture. The entire effect of user ideology on retweets from verified accounts is mediated through incoming tweets. By contrast, not only is there a minimal mediated effect of incoming tweets via unverified accounts, ideology has negligible (and non-statistically significant) overall effect on what users retweet from unverified accounts.

[47] As discussed previously, although our friends do tend to share our political views, because we pick our friends for reasons other than shared political views, homophily is not absolute (Huckfeldt & Sprague, 1987).

Figure 24*. Selective Curation Mediated by Selective Exposure*

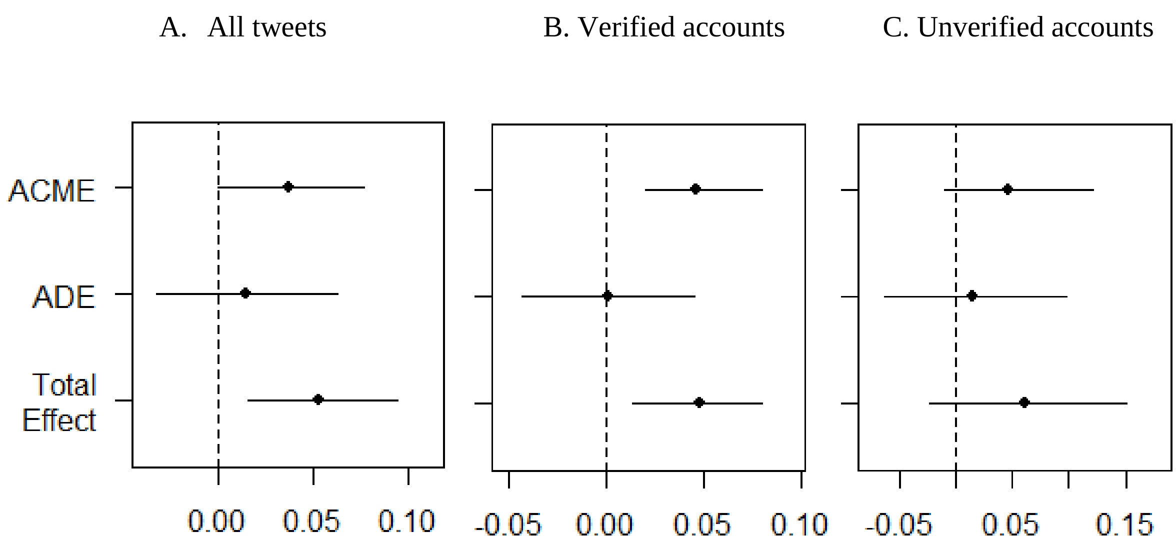

Mediated effect of users' ideology on outgoing tweets (Average Causal Mediated Effect or ACME), mediated by incoming tweets compared to Average Direct Effect (ADE). A. All incoming tweets and retweets. B. From "verified" accounts. C. From "unverified" accounts.

This last result contains some surprising – and promising – news about social media. "Social media" can be broadly used to describe platforms where any user can connect to another account for the purposes of sending or receiving information. A narrower definition of social media, however, only includes those connections that are social in nature – i.e. between two individuals who know each other as opposed to between, say, a news organization and an individual. Twitter allows for both types of connections. Connections between news organizations and users are best reflected in verified accounts and their followers; we could even say that connections to verified accounts are a variant of the traditional broadcast media model in which media institutions produce news and consumers choose which media outlet to follow. Networks of "unverified" accounts, however, represent the "social" dimension of social media (Colleoni et al, 2014).

Using this distinction, we see a clean relationship between traditional media accounts and the users that follow and retweet them: users choose to follow media that reflects their political beliefs and, in turn, share information that likewise reflects their political beliefs. When it comes to the friends we follow – the social part of social networks – however, our ideology is not so neatly aligned with the political information we receive or share with our friends. This suggests that, even with homophily, our friends expose us to information that we would not choose to consume ourselves. It also suggests that *social* media – as opposed to the traditional media model – allows users to likewise pass along a more diverse set of news to their friends. In other words, there is a weaker echo in our social media than in our traditional media use.

Analyzing Extreme Curation

In the analysis above we could not find evidence that users have an ideological bias in what they choose to retweet. But might they still have a preference for sharing *extreme* messages?

As we did in searching for ideological curation biases, we can look for an extreming bias by seeing if the proportion of extreme information that users are exposed to matches or is exceeded by what they tweet out. To do so, I bin tweets into deciles, using the Bakshy index once more, from "extremely liberal" to "extremely conservative" media sources. Examples of news sources in each decile are shown in Table 2. With these categories I can compare the distribution of tweets panelists are exposed to with what those users tweet out.

Table 2. *Examples of Online News Sites in Each Decile of the Bakshy Index*

Decile	Bakshy range	Examples
Extremely liberal	< -0.8	thenation.com, occupydemocrats.com, dailykos.com, msnbc.com
Strongly liberal	-0.8 to -0.6	newyorker.com, slate.com, vox.com, huffingtonpost.com, npr.org
Solidly liberal	-0.6 to -0.4	nytimes.com, buzzfeed.com, thedailybeast.com, bostonglobe.com
Moderately liberal	-0.4 to -0.2	time.com, economist.com, nbcnews.com, cnn.com, washingtonpost.com
Marginally liberal	-0.2 to 0	bloomberg.com, baltimoresun.com, cbsnews.com, usatoday.com
Marg. conservative	0 to 0.2	nationaljournal.com, cnbc.com, forbes.com, thehill.com
Mod. conservative	0.2 to 0.4	nypost.com, wsj.com, reason.com, examiner.com
Solidly conservative	0.4 to 0.6	armytimes.com, policeone.com, catholicnewsagency.com, infowars.com
Strongly conservative	0.6 to 0.8	christianpost.com, washingtontimes.com, newsmax.com, foxnews.com
Ext. conservative	> 0.8	dailycaller.com, nationalreview.com, breitbart.com, rushlimbaugh.com

Figure 25 shows the aggregate distribution of tweets all 472 panelists are exposed to compared to what they retweet. In the top row, we almost see a normal distribution of incoming tweets, with users mostly exposed to centrist tweets and thin tails of exposure at the extremes. In the second row, we likewise see that users mostly tweet out posts from the center, yet the distribution is considerably flattened and the tails at either extreme are far more pronounced. We see this pattern both for tweets from verified and unverified accounts, though it appears to be a stronger trend with the latter.

Figure 25. *Aggregate Extreme Curation: Panelists in Sample*

A. All tweets B. Verified tweets C. Unverified tweets

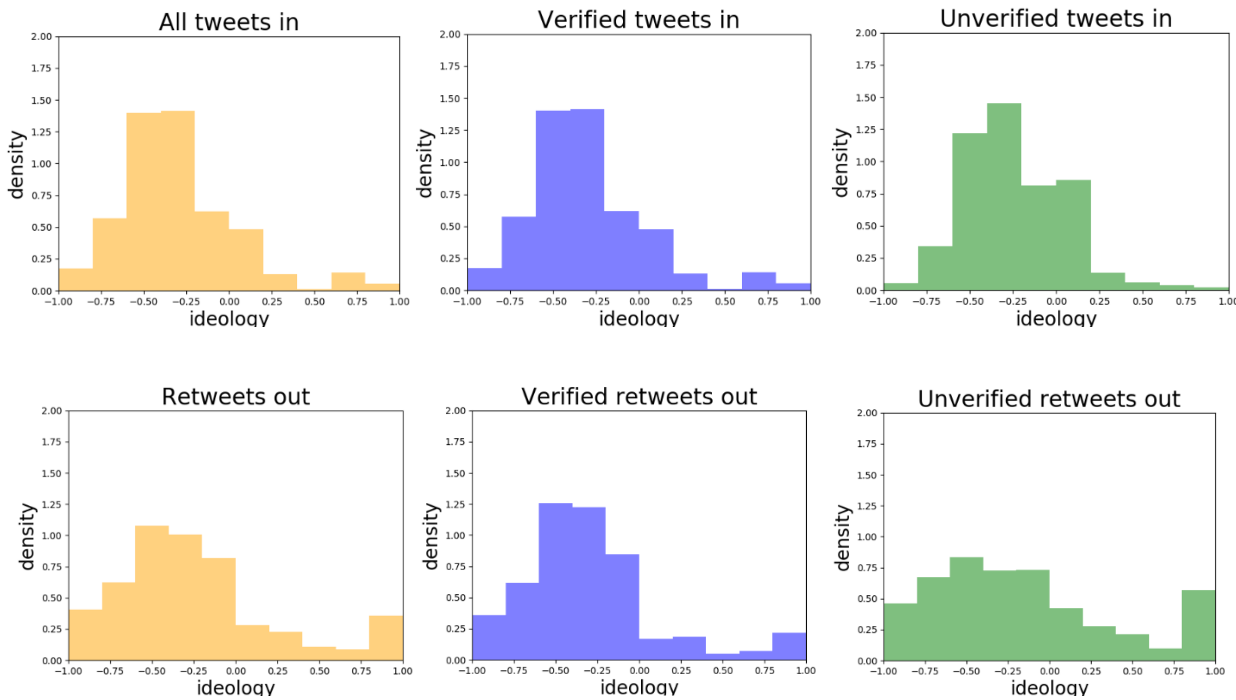

Distribution of tweets panelists are exposed to ("tweets in") and retweet ("retweets out") from the most liberal decile (-1.0 to -0.9 on Bakshy index) to the most conservative decile (0.9 to 1.0 on Bakshy index) for a) all incoming tweets, b) tweets from verified accounts and c) tweets from unverified accounts. In all cases, the distribution of articles users retweet is heavier at the tails of distribution than are the distribution of articles they are exposed to.

To check that the above results are not particular to our set of panelists, I run the same analysis with the random pool of 5,000 tweeters that were followed over the same time. We see a similar pattern in Figure 26 where, again, users retweet proportionally more extreme tweets than what comes through their feed. Also mirroring the results above, there appears to be a stronger extreming effect with unverified tweets.

Figure 26. *Aggregate Extreme Curation: Random Sample of Twitter Users*

A. All tweets B. Verified tweets C. Unverified tweets

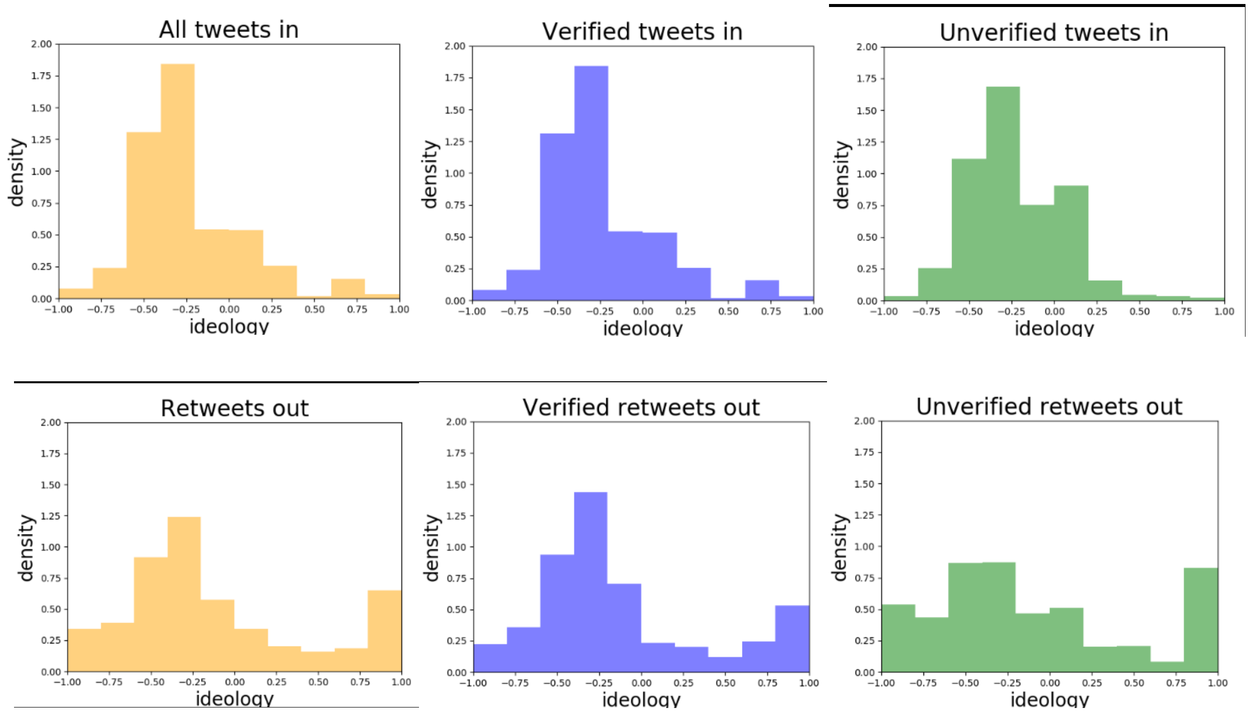

Distribution of tweets random users are exposed to (“tweets in”) and retweet (“retweets out”) from the most liberal decile (-1.0 to -0.9 on Bakshy index) to the most conservative decile (0.9 to 1.0 on Bakshy index) for a) all incoming tweets, b) tweets from verified accounts and c) tweets from unverified accounts. In all cases, the distribution of articles users retweet is heavier at the tails of distribution than are the distribution of articles they are exposed to.

These results reflect observations in other studies that look at the distribution of the political information users share. Facebook researchers, for one, chart both the distribution of unique URLs that circulate on its platform (Figure 27a) as well as the proportion of the links users shares (Figure 27b). Although there are a fair amount of “neutral” URLs in the Facebook information environment, users are not generally interested in sharing them. Instead they are strongly biased toward sharing news stories that come from either end of the media spectrum.

Figure 27. *Signs of Extreme Curation on Facebook*

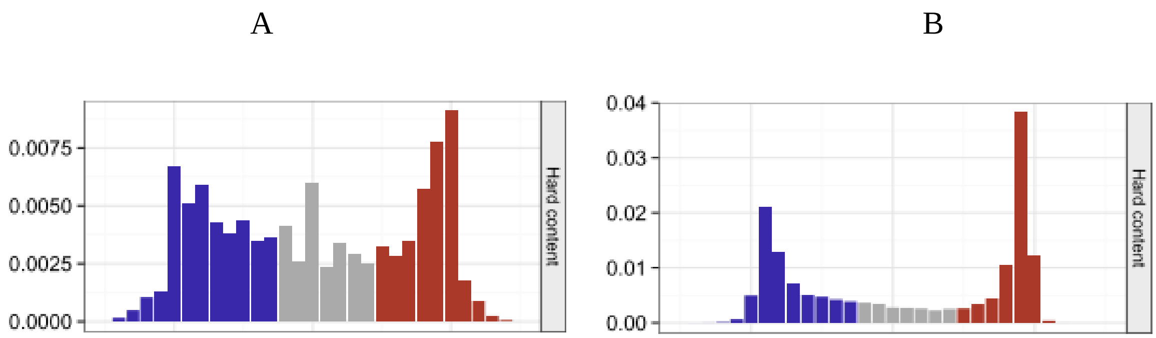

Distribution of hard news URLs shared by Facebook's self-identified liberals and conservatives, looking at (A) number of unique URLS and (B) total number of URLs shared. From Bakshy et al, 2015.

Harvard's Berkman Center researchers sees similar distributions when they looked at the news URLs users shared on Facebook and Twitter.[48] Again, as seen in Figure 28, users show a preference for sharing news stories at the farther ends of the news media spectrum.

Figure 28. *More Signs of Extreme Curation on Facebook and Twitter*

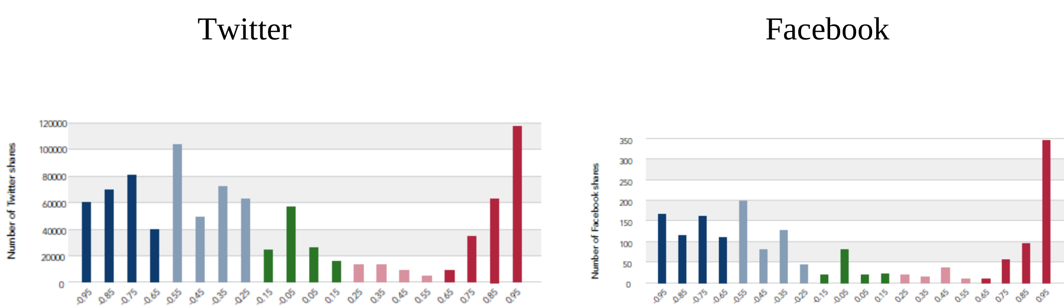

Distribution of news stories shared on Twitter and Facebook, from Faris et al, 2017.

---

48 The Berkman study used a similar method to assign ideology scores for media articles as does the Facebook team; sites' left-right score depends on the proportion of left to right users who share with those articles. A key difference with their methods is that the Berkman ideology estimates depend on identifying the ideology of users by who retweets Clinton and Trump (rather than self-identification).

The results from the Twitter data in this study – along with data from the Berkman and Facebook papers – indicate that there is an *aggregate* trend toward re-sharing extreme posts. But we still do not know whether that finding driven by the extreme tweeting of a small number of avid users – or is it a broad tendency across users? To see, I compare the average proportion of moderate (quintile 1) to extreme (quintile 5) tweets that individual users are exposed to and that they retweet, with results seen in Figure 29. For both verified and unverified tweets, users on average tend to retweet out fewer moderate tweets (1) and proportionately more extreme tweets (quintiles 4 and 5). As we saw above, the bias toward retweeting extreme tweets is even starker among unverified tweets. Again, checking that this pattern is not particular to the panelists, we see similar results among the random pool of 5,000 Tweeters.

Finally, one might be curious to see if this a bias toward re-sharing extreme tweets on one's side of the ideological divide – or a general propensity to retweet extreme stories. There are at least two reasons why a liberal user, for example, may want to share an extreme conservative post; they may want to mock that post or to alert her friends to the "outrageous" information the other side is spreading. To get a sense of whether an extreming bias is one-sided, I look at only the liberal panelists and the liberal tweets they see and retweet. As can be seen in Figure 28c, the extreming pattern is diminished, but is still visible among unverified tweets. Due in part to the smaller sample, N=89, it is also not a statistically significant effect. More analysis is needed to see of this diminished effect among verified tweets suggests that at least part of any extreming bias includes a propensity to cross-share extreme tweets – or if the results for all panelists are driven in large part by a small subset of conservative tweeters.

***Figure 29**. Average Extreme Curation*

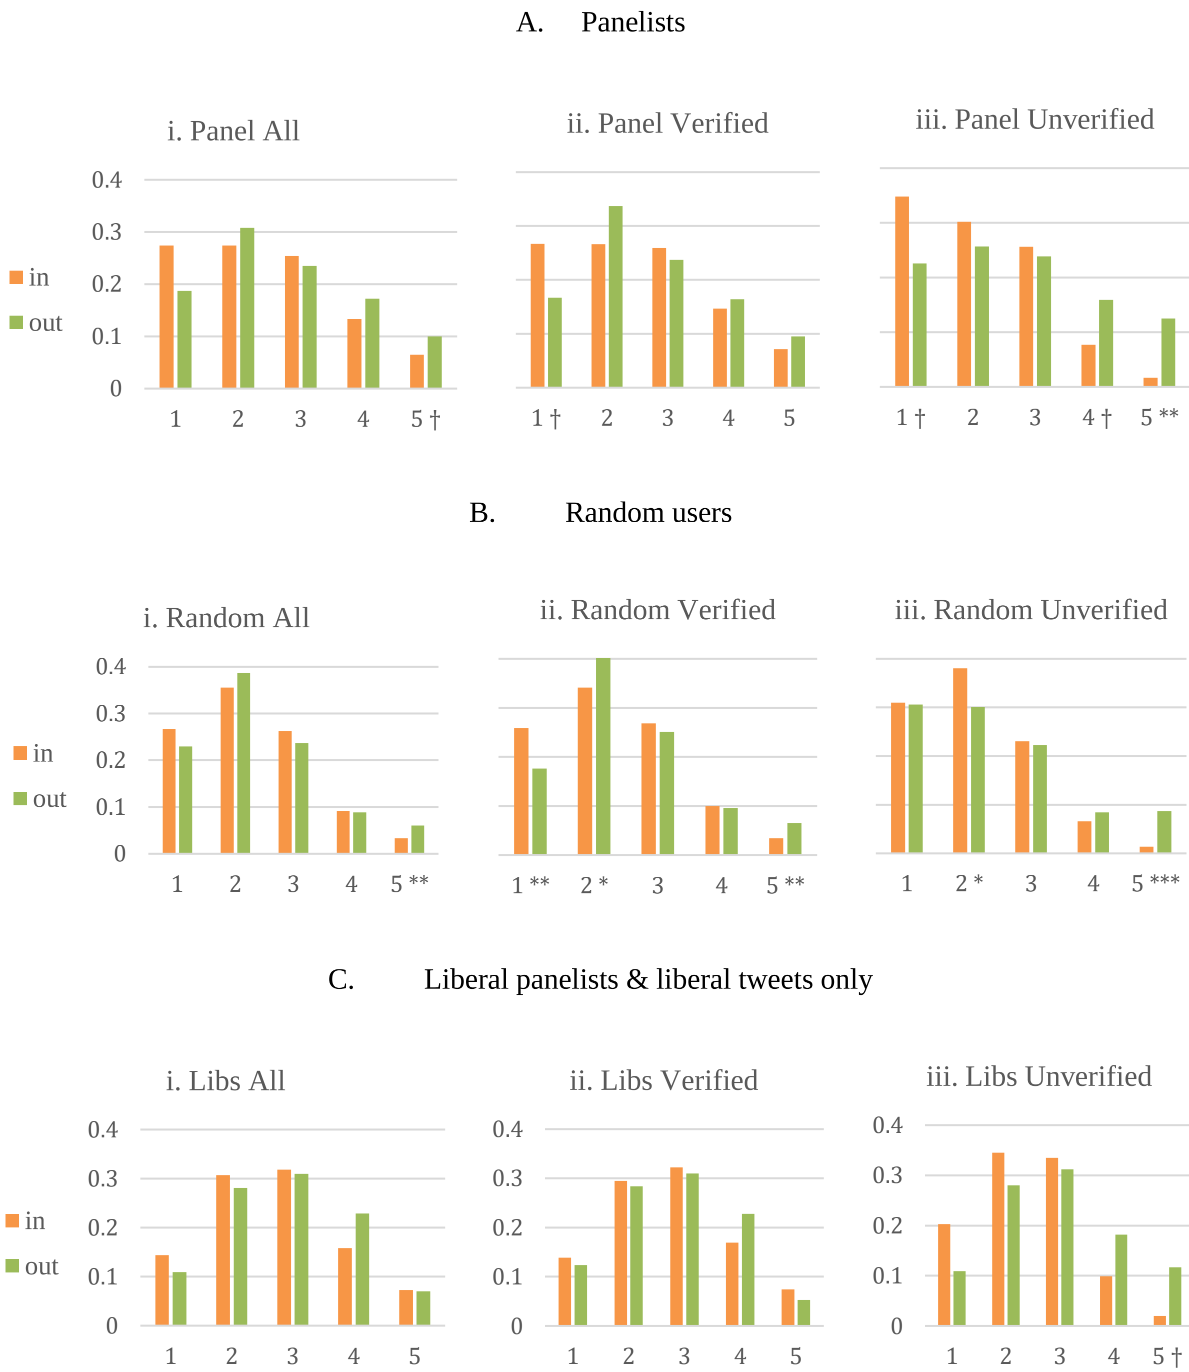

Proportion of tweets users are exposed to ("in") and retweet ("out") that are moderate (1) to extreme (5) for: (i) all tweets, (ii) tweets from verified accounts and (iii) tweets from unverified accounts. (p-values: †<0.1, *<0.05, **< 0.01, ***<0.001)

What explains greater extreme curation of tweets from unverified accounts than from verified accounts? Two factors may combine to make extreme tweets from our friends so re-sharable. First, our friends are best suited for selecting *any* news story that we find eye-grabbing and worthy of sharing – for they themselves were drawn to share those stories and, as our friends, they have similar sensibilities. In general in this sample, posts from unverified accounts are retweeted at a rate 30 times higher than tweets from verified accounts.[49] Also, as posited in this dissertation, extreme tweets have qualities that make them exceptionally retweetable – being sensational, alarmist, emotionally-valenced, etc. If our friends are pre-selecting the extreme tweets that they find particularly tweet-worthy, we may be seeing a powerful interaction effect between the inherent tweetable-ness of extreme posts and our friends' ability to select information that is in sync with our preferences.

The causes behind users' predilection to re-share extreme posts will be explored in Part II of this dissertation. At this point, we can fairly confidently say that Twitter users are, indeed, extreme curators. The impact on the distribution of information on social media could be substantial; to the extent that users re-share political information, they will shift the tenor of the news environment online from moderate to extreme. In this chapter we observed only one wave of re-sharing; as we saw in Chapter 2, when we compound users' preference for re-sharing extreme information over subsequent waves of diffusion, the dominance of those extreme tweets only increase.

Summing Up our Curating Behavior

While research – and intuition – tell us that social media users prefer to post news that align with their political leanings, no previous research examines if users are indeed ideologically biased in the political information they choose to share. To show such a curation bias, it is necessary to see

[49] This could be explained in large part by the limitations of data collection which likely result in a fuller sample of incoming tweets from unverified than verified accounts.

not only what users post, but the set of information they select from. In running such an analysis using the data of 472 Twitter users, evidence presented in this chapter suggests that users may in fact be ideologically *un*biased in what they retweet; instead of cherry-picking ideologically congruent news stories to retweet, they appear to retweet a representative sample of the tweets that come through their feed.

The data does, however, confirm previous studies showing that users are de facto ideological curators. These combined findings suggest that whatever ideological bias exists on social media occurs at the level of selective exposure (the accounts users choose to follow) rather than selective sharing (which tweets they choose to repost). Mediation analysis in this chapter confirms such a relationship.

It is important to note, however, that the results in this analysis are null findings; the lack of an observed curation bias does not mean one does not exist; rather it may be a consequence of a small sample size. To confirm the absence of an ideological curation bias, in would be necessary to analyze a larger sample.

At the same time, while the data presented here do not indicate that users have an ideological curation bias, the data do show a pronounced bias toward sharing extreme information. Such a preference appears not only among the same 472 panelists analyzed above, but also among 5,000 random Twitter users followed in this study. This pattern suggests that, while an ideological curation bias may exist if we were to examine a larger data set, it would likely be dwarfed by users' bent toward re-sharing extreme information. We are not partisan curators as much as pushers of extreme information.

Chapter 5. Re-parametrizing the Model

In the last three chapters, I isolated and examined each of the three mechanisms that dictate how information is distributed in a social network. First, we modeled how ideological and extreme information flow through homophilous networks, assuming varying levels of homophily and biased curation on the part of users. We saw that diffusion's capacity to increase information sorting and extreming is a complex interaction that depends upon the level of homophily in the network as well as users' biases in what they post and re-post. We then looked at what those levels of homophily and biased curation might be, first examining estimates for homophily on two existing online social networks, Twitter and Facebook, and then analyzing curation behavior among a sample of Twitter users.

We now come full circle. Given estimates of homophily and curation that we see in real social networks, what can we say about social media's capacity to create information bubbles or propagate extreme information? To grasp what is happening inside the "black box" of social media, we will want to re-run the models in Chapter 2 using the parameters identified.

First, consider social media's ability to sort congruent information and create information bubbles. The parameters we identified in the sections above are a level of homophily of approximately 0.7 and a re-posting bias of 0.5 (that is, no discernable bias). To be consistent with the models presented earlier, I stick with networks of 1,000 nodes, in which nodes have an average of 8 connections, and I run 200 simulations on each graph generated.

In setting up the simulations, I modify one additional parameter. In chapter 2, recall, I used a 10% ceiling rate of contagion, chosen for the practical reasons that rates higher than 10% risked leading to messages saturating the network and rates lower than 10% too often saw messages die out quickly (i.e. rarely reaching a 5^{th} wave of diffusion). This rate, however, is far higher than what we know about actual retweet rates on Twitter. The full sample of 5,471 Twitter users I observe, for example, retweet about 3.3% of the unverified tweets that come through their feed. To better reflect real diffusion patterns, I use a ceiling rate of 6% which results approximately 3.3% of messages being reshared. To compensate for the high frequency of messages dying out, I run simulations on 1,000 instead of 100 generated networks so we are assured of observing cascades that run for five waves.

This ceiling rate of contagion has an added benefit for our model as it produces a ratio of initial posts to re-posts that reflects real ratios on social media. This ratio is relevant because ultimately we will want to know not just how much congruent and extreme information users are exposed to in each wave of diffusion, but how much congruent and extreme information users are exposed to in total *across* all waves. That ratio in social media is hard to pin down, but one estimate (Liu et al, 2014), puts the proportion of posts on Twitter that are retweets at 25%. Using a ceiling rate of contagion of 6% results in approximately 25% of messages that nodes observe being re-shares.

Finally, we need a parameter for nodes' initial posting biases. This number is the likelihood a node posts an initial message that reflects their ideology. Again, this is not a well-known parameter. For the purposes of this simulation, I use the average proportion of original congruent tweets among the liberal Twitter users in my sample who posted at least two original tweets; that proportion is approximately 70%.

Plugging those parameters into our model (using a small world binary graph), we can see the effects of diffusion in Figure 30.

Figure 30. *Modeling Information Sorting with Real Life Parameters*

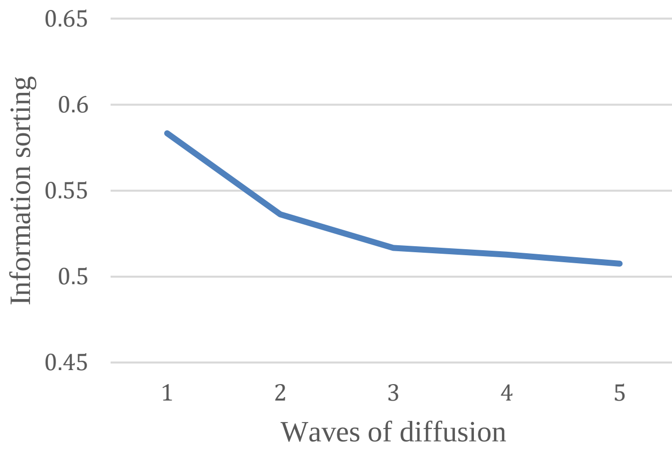

Levels of information sorting across waves of diffusion in agent-based model, using estimated parameters from real online social networks: 0.7 average homophily, 0.5 ideological curation and an initial ideological curation bias of 0.7.

Figure 30 presents only the weakest of information bubbles – and one that diminishes with greater diffusion. At the first wave of diffusion - that is, users' initial posts - the average user will be exposed to 58.4% congruent information. The more a news article is re-shared, however, the more evenly it will be distributed between ideologues and, likewise, the more likely a counter-ideologue will see it. Put another way, the re-tweets that users see will be more ideologically balanced than the original tweets they are exposed to. If an article were to be re-shared 3 times sequentially (a rarity even among viral tweets), by the 4th wave we might expect it to fully break down any information bubble, being evenly distributed between left and right users.

To assess social media's capacity to fuel the spread of extreme information, I run the model above with additional modifications. As above, I run simulations on a binary small world network with a homophily rate of 0.7. This time, however, I chart the spread of two types of information; "moderate" information for which users have a curation bias of 0.5 (i.e. no bias) and "extreme" information for which users have a 0.9 curation bias. While these levels of bias may seem far apart, in the model they mean that a user is only 9/5 times more likely to re-share an extreme (congruent) post than a moderate one; our Twitter analysis above, however, suggests that the relative re-shareability of extreme tweets is even greater, so using these parameters will likely underestimate any gulf between the viral spread of moderate compared to extreme posts.

With those parameters, liberal and conservative users will share moderate messages with equal frequency (0.5 * Ceiling rate) while they will have an inverse sharing bias for extreme messages (0.9*Ceiling rate for congruent and 0.1*Ceiling rate for incongruent).[50] To further reflect what we see in the Twitter data in the last section, users in the model are themselves initially exposed to little extreme information - just 5% of the information they can select from will be extreme and ideologically congruent, while 90% will be neutral and 5% is extreme information from the other side. We can think of these initial posts as the information they get from "verified" accounts or from outside media sources. As in the previous simulation, users have an initial curation bias of 0.7.

Figure 31. *Modeling Extreme Curation with Real Life Parameters*

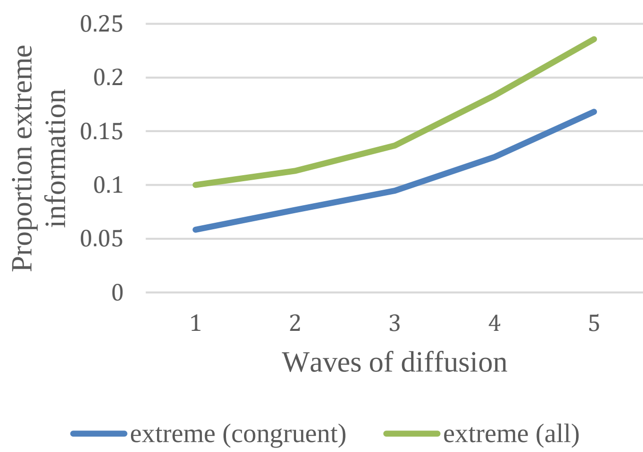

Proportion of messages that are extreme that users are exposed to at each wave of diffusion in a network with 0.7 homophily, assuming that extreme messages are 10% of all initial posts in a network (5% liberal and 5% conservative).

---

[50] We are assuming, again, that any extreme bias is purely for stories on one's side of the spectrum, but as discussed above our data does not tell us if the bias we see is for *any* extreme story, regardless if it is left or right. If an extreming bias is indifferent to congruence, then the effect we see in Figure 31 would be stronger.

Figure 31 depicts users' exposure to all extreme information as well as to extreme information that is congruent with their beliefs. In contrast to information sorting, Figure 31 shows that diffusion fuels information extreming. Without diffusion, users are exposed to relatively little extreme information, but as messages are re-shared, extreme messages make up an increasingly large proportion of what users are exposed to. If we look down the path of diffusion for 5 waves, the proportion of congruent extreme information users are exposed to will have more than doubled.

The figures above, however, only show us how information sorting and extreming fluctuates in each subsequent wave of diffusion. Ultimately, we want to know what the *cumulative* effect of diffusion is on sorting and extreming; when we add up the messages that nodes are exposed to across all waves, what proportion are congruent and extreme? Because only a small number of messages are re-shared, nodes see relatively few re-shares overall; we should thus expect an overall modest effect of diffusion. That is what we see. Adding up all the messages nodes are exposed to across the five waves of diffusion, information sorting lands at 57%, just a 2.4% decrease from initial sorting levels. Turning to information extreming, the cumulative effect of diffusion leaves nodes exposed to 10.7% extreme information, a 7% increase from the 10% of extreme information our model exposed them to before diffusion.

All told, then, we can say that the extended network users are embedded in and political information's ability to diffuse through that network have negligible impact on information sorting; what effect there is on sorting, moreover, will be numerically negative - or positive, from the normative perspective of those who push for exposure to diverse information. Diffusion's effect on extreming, however, is both marginally greater and will be numerically positive - or negative for those who worry about polarization.

As stated at the outset, the dominant factor in determining how closed is one's echo chamber is one's selection of friends. It is that first tier of friends and their initial posting biases (i.e. what they pluck from their outside news environment) that predominantly decides how homogenous one's social media information environment is. The diffusion of information through the network – the resharing of like-minded stories that is feared to amplify our bubbles – has little effect on the homogeneity of our social media feeds, and what little effect exists opens us up to more challenging information. All told we can expect, knowing what we do about users' choice of friends and their sharing and re-sharing biases that they will be exposed, on average, to about 43% incongruent information.

That number aligns somewhat with what Facebook researchers find in their analysis of users' exposure to cross-ideological news (Bakshy et al, 2015). As seen in Figure 32, Facebook reports that conservatives are exposed to about 35% cross-cutting information while liberals in their sample are exposed to approximately 24% incongruent information. Those numbers are lower than the 43% predicted above, but a key distinction in the Facebook analysis is that it does not include "moderate" information that falls on users' side of the political spectrum; if moderate left and right information were to be included we would expect the study's estimates to rise.

Figure 32. *Information Sorting on Facebook*

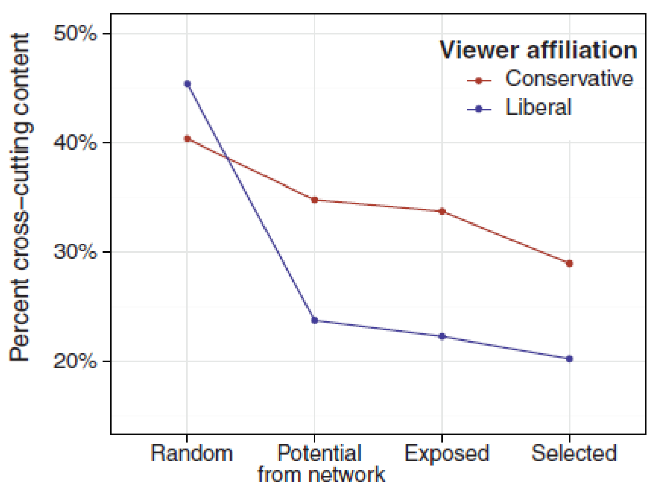

Exposure to cross-ideological news stories by self-identified ideologues on Facebook. From Bakshy et al, 2015.

While bursting the bubble on social media's complicity in creating echo chambers, the analysis in Part I gives us cause to believe that social media does contribute to the spread of extreme news. The effect may be small; again, our models suggest that users' biases in resharing information from the far left and right could lead to about 7% greater exposure to extreme information. But given the possible impact on users' perception of their political opponents, particularly if we imagine that the most viral stories are the most incendiary and salient, then even a marginal increase in exposure should be troubling.

Again, though, the story may not be about how much social media amplifies extreme information via diffusion but about the kinds of stories users introduce into social media environments in the first place. As suggested by the Facebook and Berkman study graphs we saw above, Facebook and Twitter users have a distinct preference to share (whether re-sharing or introducing into their social media environments) stories from the ends of the political spectrum (even more so than the Twitter users in our sample). It is that mix of news that makes a potentially polarizing brew of information on social media.

Why social media users have this apparent bent toward sharing news from the ends of the political spectrum is one of the central questions of the second half of this dissertation. In it I construct a theory to explain what motivates social media users to share any political information and why those motivations would lead users to share not only congruent, but "extreme" news stories. In understanding those motivations, my hope is that we might not just arrive at greater knowledge but also better insight into how to mitigate the toxicity of social media's news feeds.

Part II. A Theory of Information Sharing: Why Social Media Users are Motivated to Share Congruent and “Extreme” Information

Chapter 6: Building a Theory of Political Information Sharing

There is nothing preordained about what counts as "news." Out of the infinite shapeless and complex events, trends and stories that happen around the globe on a given day, only a relative handful of occurrences end up as a story in print or on the air. Which events become news stories is decided by media's "gatekeepers" – the writers, editors and media owners who decide which stories to produce. In Part I, we were introduced to one such gatekeeper, "Mr. Gates," the wire editor of a small Midwestern newspaper, whose gatekeeping decisions were studied by David Manning White. While White's study was groundbreaking methodologically, it did not reveal strong findings about what makes a story worthy of passing a gate (it hardly could, having studied one editor's choices for one week); it did, however, set in motion decades of research aiming to understand news gatekeepers and the motivations behind their decisions. Thanks to subsequent scholars we now know, for example, that a host of biases, both personal as well as widely held "American" values, shape the stories editors choose (Gans, 1979). We also know that practical constraints dictate what gets reported; since news companies depend on audience eyeballs and advertiser money, editors' choices are a calculus of how inherently "news-worthy" a story is, how much it costs to produce, how it might increase readership and how it could potentially aggravate advertisers (Shoemaker et al, 2001: Gans, 1979; Soroka, 2012; Sigal, 1973).

In the past, those gatekeepers decided not only what stories to run, but which stories consumers were more likely to see. Whether on the pages of the New York Times, on the evening news or on news radio, gatekeepers curated the order and prominence of those stories. In print media, readers are free to choose what articles to read by scanning and skipping stories, but editors direct readers to the stories they deem most important by placing those stories "above the fold." On TV and radio, consumers are entirely captive to the curation decisions of the news editors.

With the internet and social media entering the media landscape, news institutions still maintain the dominant gatekeeper role in deciding what news stories are produced, but other players - bloggers, search engines and, now, lay citizens - increasingly decide what news we see. In 2017, 67% of Americans reported getting at least some of their news from social media, up from 62% in 2016 (Shearer and Gottfried, 2017). According to one media consulting group, as much as 43% of its clients' news content was reached via social media (Ingram, 2015). While news institutions have a presence on social media, particularly on sites like Twitter, and thus still curate much of the news, it is one's "friends" on these sites who increasingly decide which stories come across our laptop and smartphone screens. News producers largely determine which stories get told and how, but our friends have become our new news curators.[51]

Just as scholars once probed what motivates the gatekeeping decisions of news producers - the biases, constraints and incentives that drive their choices - today, if we want to understand the information that news consumers see, it is critical to understand what drives social media users to post the news stories they do.

In the previous sections we saw that those choices matter; they determine how homogenous our ideological bubbles are and how extreme the news is that fills our feeds. The diffusion models presented in Part I showed us that, to the extent that users have ideological biases in what they share, they determine how strong their friends' echo chambers will be (although they cannot make those echo chambers more homogenous than their friends' networks). If users are motivated to share extreme information, moreover, they can collectively amplify extreme stories, making views from the far left and right dominate their friends' information environments.

---

[51] Although not discussed in this dissertation, many others have made the case that social media also increasingly influence which stories traditional media report (Zuckerman, 2018).

In this part of this dissertation, I shift from *observing* our online sharing behavior to trying to *explain* it. I speculate on the forces that drive these new gatekeepers to share political stories and ask how those motivations explain why users share the stories they do. Specifically, I aim to understand why we tend to share articles that reflect our ideology and that are extreme in nature.[52]

I start by considering whether an existing theory – "selective exposure" – can answer those questions as other scholars have proposed. I conclude that selective exposure theory has limitations in its explanatory power for "selective curation." I present an alternative framework for understanding social media users' curation choices that builds on theories in the fields of Social Psychology, Sociology, Anthropology and Communications. I propose that when users post on social media, they are broadcasting to an "imagined audience" that represents one of their social groups (Ellison & Vitak, 2015). In communicating with those conjured groups, users are driven by two motivating forces, what I term "group impression management" and "group rallying." "Group impression management," which draws from Impression Management theory and Social Identity theory, posits that when posting political information users are primarily motivated to project an image of themselves that will bolster their inclusion and status in those groups (Goffman, 1959; Leary et al, 2015; Tajfel, 1974). That image will be of one who is a loyal and competent group member. As a secondary force, "group rallying" motivates users in times of threat to seek support and coordinate group action by drawing members' attention to threats to the group. Both those sets of motivations lead users to selectively post political stories that confirm the group's beliefs and superiority, and thus are congruent with their political affiliation, as well as stories that are "extreme" in nature – i.e. that are dogmatic, emotionally-valenced and tribal.

[52] The reader may recall that in Chapter 4, we found no ideological curation bias among Twitter users, yet I conjectured that such an absence of bias might be explained by the biased news environment they create for themselves. We might still suppose that users would exhibit a bias if presented with an equal set of liberal and conservative news stories. In the next chapter, we will see that is the case.

In the chapter that follows this one, I test elements of that framework in a series of online survey studies. First I test whether social media users prefer to share stories that confirm their political group's views and cheer the ingroup. Recall that in Part I we did not find such a curation bias among Twitter users, but I conjectured that was the case because Twitter users had created an information environment that obviated the need to be ideologically selective in the stories they shared. In the next chapter, I create an environment where I can discern a preference for sharing stories that are politically congenial. I use such an environment to also test to see if users' have a preference for sharing dogmatic stories and stories that explicitly praise the ingroup or denigrate the outgroup. Next I prime the salience of subjects' social groups in order to detect signs that they are, indeed, communicating with one of their social groups in mind. Finally, I manipulate subjects' need to be included to see if I can likewise increase their motivation to project an image of themselves as loyal group members. I find support, albeit sometimes weak, for each of the propositions.

Understanding the motivations behind users' choices on social media is not just academic. While as scholars we care about what drives political behavior, having a window into the psychological mechanisms behind users' choices gives us insights into how to temper their more destructive choices. Social media platforms, in particular, have an interest in mitigating the spread of stories that may spark inter-group violence; understanding what propels users to share those stories gives us a handle on how to deter possible harmful instincts to click and share.

Selective Curation: A Mirror of Selective Exposure?

Before proposing a new theory to explain any behavior, one should first see if an existing theory can adequately do the job. In this instance we need to first ask: is there an established theory that can explain what drives social media users' curation of political information?

Joseph Cappella, Allison Earl and their colleagues would say there is such a theory – "active selective exposure" (Cappella et al, 2015; Earl et al, working paper). Selective exposure, as discussed in Part I, predicts that citizens tend to consume news aligned with their political beliefs rather than challenge themselves by reading news from the other perspectives. “De facto” selective exposure can be explained in large part by environmental factors; the friends we have and neighborhoods we live in surround us with news that aligns with our beliefs. But studies show that we are also *actively* drawn to information that confirms our beliefs and avoid information that may challenge those beliefs (Frey, 1986; Hart et al, 2009; Stroud, 2008).

The underlying mechanism that drives active selective exposure is, according to most studies, Festinger’s theory of cognitive dissonance. That theory is often understood in the simple way Festinger put in it in 1962: the “idea that if a person knows various things that are not psychologically consistent with one another, he will, in a variety of ways, try to make them more consistent” (Festinger, 1962). Dissonance is an uncomfortable feeling that people try to mitigate; one way to do so is to avoid dissonant information in the first place. It follows that if someone holds belief A, they will eschew belief anti-A to avoid discomfort. Yet Festinger’s foundational experiments and subsequent research by other scholars show that dissonance is more complex than a simple mechanism that abhors inconsistency. Underlying dissonance is not a discomfort with holding two contradictory beliefs, but rather with being faced with information that challenges one’s sense of self or previous decisions.

The central role of the self can be seen in one of Festinger’s early experiments and a variation of that experiment conducted by Claude Steele. In Festinger’s original study (Festinger & Carlsmith, 1959), he and coauthor Carlsmith asked subjects to lie to other participants (who were confederates in on the ruse) about an experiment they just completed. When subjects lied with little incentive to do so - merely being paid an extra $1 - they experienced a dissonance between what they believed and what they publicly said. To relieve that dissonance, according to Festinger, subjects adjusted their beliefs to be more in line with their stated position. Another group of subjects who received $20 to lie, however, did not likewise experience dissonance and consequently adjust their beliefs. Unlike the group that received $1, Festinger explains, those subjects could tell themselves that they were essentially forced to lie ($20 is the equivalent of $180 today) and so could convince themselves that there really wasn’t an inconsistency in their beliefs and behavior.

Claude Steele suspected something more might be going on for the group that received $20. Steele tells us that humans have a need to maintain a view of themselves as "adaptively and morally adequate" (Steele, 1988). We can get that sense of adequacy either by seeing ourselves as competent and strong or as good and morally upright. If one's sense of self takes a hit - for example by being inconsistent (and so either incompetent or duplicitous) - one will try to restore a positive sense of self. For the subjects in Festinger's $1 condition, Steele would argue, the only way to do so was to remove the inconsistency, adjusting one's beliefs to be in line with one's behavior. But, according to Steele, it is not necessary to restore a sense of self *in the same domain* that was diminished. In the face of evidence that one is inconsistent, one can restore a sense of self by searching for other evidence of one's value. That is what subjects in the $20 condition were able to do; they could reflect on how smart they were for scoring $20 (again, $180 in today's dollars) for telling one measly harmless lie.

Steele tested this hypothesis by running a version of Festinger's experiment, but adding a "self-affirmation" manipulation for a random set of the subjects, giving those subjects time to reflect on an event that made them feel strong in an area they deemed important (family, career, morality, etc.). When subjects had an opportunity to "affirm" their sense of self, they had less of a need to resolve any dissonance; it no longer bothered them that they lied since whatever deficit their inconsistency created in their sense of self was restored in another domain (Steele, 1988).[53]

Steele's work sheds light on the kinds of information we will seek to consume or avoid. For one, we will want to avoid information that suggests we have "knowingly chosen to engage in a bad or foolish behavior," because that will be "inconsistent with a self-image as a decent and intelligent person" (Kunda, 1990). As one example, offered by Frey in his review of selective exposure (1986), a new homeowner will seek out information that confirms he made the right purchase - while avoiding information that shows he got a bad deal - to avoid feeling like he's

[53] A tendency to seek information that is consistent with our current beliefs – or "confirmation bias" – may also be explained by factors other than a psychological drive to perceive ourselves as competent and good. When asked to seek information to see if a theory is true, humans naturally look for information that confirms rather than disconfirm the theory (Klayman, 1995) .This logical lapse, which runs counter to scientific inquiry, seems to be a simple cognitive quirk, not a psychologically driven bias.

“been had.” Likewise, if we love a good hamburger we might avoid information that claims that eating red meat is unhealthy, unethical or destructive to the environment.

But it is not just information that is inconsistent with our *behavior* that we will seek to avoid. We will also avoid information that challenges our beliefs and values, particularly those we hold dearly or wear publicly on our sleeves. At one level, we avoid information that challenges our beliefs for the same reason as the homeowner above; information that reveals one is “wrong” impugns our competence by making us look stupid.[54] At a deeper level, however, information that challenges our values and closely held beliefs will be a threat to our identity (Stitka, 2002). Those deeply held beliefs and values, like our behavior, contribute to the image of ourselves as good and competent human beings. Our values ground our lives in a narrative in which we are a sympathetic protagonist - whether that is a responsible homemaker, a fighter for justice, or a champion of capitalism (Lane, 1962). Our beliefs are likewise necessary to maintain that personal narrative; to be a fighter for justice you must believe injustices exist, and one must believe that capitalism is good to be its champion. A challenge to those values and beliefs is then a challenge to *who we are* and so should be avoided at all costs.

One part of our identity - our social identity - plays a particularly prominent role in the realm of politics and the information we choose to expose ourselves to. Social identity theory claims humans have a special proclivity to associate with and identify with groups (Tajfel, 1974). When we do so we take on the values and beliefs of those groups and seek to protect those values and beliefs as our own (Suhay, 2015). Thus any new story that challenges the sanctity or competence of our group will likewise be sidelined in favor of stories that confirm our group is virtuous and strong (Green et al, 2004; Taber & Lodge, 2006).

Despite our motivation to preserve our beliefs, values and identity, active selective exposure is not absolute (Kunda, 1990). Alongside our affinity for confirmatory information, humans also seek out new and challenging information. Indeed, it’s hard to imagine a cognitive species

[54] Studies show that the more publicly we hold beliefs, the more committed we are to a belief and the more a belief is connected to our values, the more we resist finding out we’re wrong (Hart et al, 2009).

surviving if that were not the case; to navigate our environments and make choices that benefit our well-being, we cannot be blind to reality. In seeking information, thus, humans have dual goals - to preserve their beliefs and to form an accurate view of the world.

Whether an individual seeks new or confirming information depends upon which goal is dominant. When information is not a threat to an individual's self-image - that is, when they have not committed to a belief or do not have values connected to a belief - they can freely seek new information. Even if we have strong beliefs, when our survival comes into question or there is a sizable tangible gain at hand, we will be willing to seek and consider information that challenges those beliefs (Bullock et al, 2013; Festinger, 1962; Kunda, 1990). Likewise, we seek out attitudinally dissonant information when we experience anxiety, which sends a signal that there are risks in our environment (Valentino et al, 2008).

In the realm of politics, however, accuracy goals will rarely trump the motivation to preserve one's beliefs, for the simple reason that there is rarely a large payoff for holding accurate information about policies or candidates. For the average voter, one's political beliefs have little material bearing on one's well-being; our single vote at the poll booth is neither likely to swing an election nor result in different policies (Riker and Ordeshook, 1968). Without significant benefits to holding accurate political beliefs, we might as well keep the ones that maintain our positive self-image.

But there is another reason active selective exposure is not ironclad; even if we are motivated to preserve our beliefs, we have other defenses against challenging information. Selective exposure is only one way we act as "motivated reasoners" seeking and processing information to reach a desired conclusion. If we expose ourselves to an article that confronts our political views, we tend to home in on the evidence that confirms our views while ignoring any challenging information (Lord et al, 1979). We can also be adept at refuting any disconfirming information if we are equipped with facts and arguments. That is why the most policy savvy are willing to dabble in articles written by the opposition; with the skills to debunk opposing information they are immune to dissonant information and can easily keep their worldview intact (Frey, 1986).

In sum, selective exposure theory explains that humans will tend to seek information that bolsters their self-image as good and competent individuals, particularly beliefs and values that are at the core of their identity and that are held by the groups they strongly identify with. At the same time active selective exposure is not absolute; humans are also partly motivated to seek out information that gives an accurate view of the world, and - when confident in their ability to refute challenging information - willingly expose themselves to opposing views.

Limitations of Selective Exposure Theory to Explain Selective Sharing

Capella and colleagues (2015) argue that the same motivations that drive selective exposure will be in force when we choose information to share with our friends: we will likewise be drawn to share information that reinforces our friends' views, yet at the same time be motivated by the - at times - conflicting motivation to share accurate information

Capella does not tell us *why* we would expect the same psychological forces behind our choices on what information to consume and what to share, but Earl and colleagues (2018) provide a clue. While not giving an explicit reason, they use the phrase "vicarious selective exposure" implying that users are practicing a form of empathy, mirroring the emotions and desires of others and so acting as if they were them (Wondra & Ellsworth, 2015). Earl provides experimental evidence that we do indeed practice "vicarious selective exposure"; when asked to select information to share with others, subjects tend to select information that is congenial to the receiver's views while withholding uncongenial information. It may be, then, that when it comes to consuming and sharing information, we simply do unto others as we would have done unto ourselves.

Vicarious selective exposure runs into trouble as an explanation for sharing political news, however, when we consider contested and divisive information. For one, when the recipient of information has views that conflict with the sharer there will be a conflict between selective exposure and vicarious selective exposure; the information that will be congruent with the recipient will be incongruent with the sender. Which will dominate, the sharer's desire to protect her beliefs or her empathetic desire to protect the recipient's beliefs? If the reader has ever witnessed a political discussion between a conservative and a liberal the answer may seem self-evident; not only will each discussant not try to preserve their ideological opponent's beliefs, they will likely try their best to dispossess their opponent of their ludicrous views. Vicarious motivations, thus, may only exist when the recipient shares our views.

Yet even in the case of recipients who think like us, we may not share the same information we are willing to consume ourselves because we may not trust the recipient to have the same powers of (motivated) reasoning. Recall subjects who are confident in their ability to refute contradictory information are more willing to read challenging news articles; their debunking skills act as inoculation against incongruent information. But when we share with others, we may not be equally confident in their skills to debunk. Whereas we are certain we are able to see the truth, humans tend to think others are susceptible to bias and manipulation (Pronin et al, 2002; Ward et al, 1997). To protect those gullible others, then, we may be more careful to sift out news that might threaten our shared beliefs.

In brief, there are limitations to how much vicarious selective exposure can explain our motivations for sharing political information; it can only do so when the sender and receiver share a worldview, and even then a sharer, in making a common assumption about others, may attempt to protect the receiver from views that we do not share.[55]

Indeed, both Earl and Capella concede that selective exposure theory may not fully explain what drives our decisions to share the information we do. They both propose an additional motivation: to strengthen social ties. As Capella says, the selection of what to share "can facilitate or hinder social goals" and there are some stories that users would read but not share. As evidence that

[55] At the same time there is at least one reason we might expect selective exposure and selective sharing to appear similar, if only because it would be hard to explain why a user would share something they had not (willingly) consumed themselves first.

social goals are in play, Earl finds that vicarious selective exposure is stronger when subjects like the person who receives the information.[56]

I agree that users are motivated by social goals when they share political information, but whereas Capella and Earl suggest such goals may be peripheral or secondary, I propose that social goals are central and primary. As such, understanding these social motivations will take us much farther in explaining users' sharing behavior.[57]

---

[56] Also worth noting is that Earl's work is dyadic, with subjects sending information to one user. As will be discussed at length later, we should expect motivations for sharing information when "narrowcasting" to a friend or two to be different from when we "broadcast" to a large audience (Barasch & Berger, 2014).

[57] While the study of information sharing is relatively nascent among political scientists, other fields have studied information sharing more extensively. Information and management scholars, for example, have examined what motivates "knowledge sharing" in organizations' online communities (Chiu et al, 2006; Hsu et al, 2007; Lin, 2007; Kankanhalli et al, 2005). Researchers on social media have likewise studied the factors that lead to personal "disclosure" - the sharing of both emotional and functional information - online (Bazarova & Choi, 2014; Rime, 2009).

One of the most comprehensive attempts to understand the motivations of information sharing, though, comes from marketing researcher Jonah Berger. Berger (2014) reviews a diverse literature and categorizes those motivations into five main functions: impression management (sharing information that creates a favorable impression of ourselves to others), emotion regulation (seeking others' input when we experience emotional confusion or distress), information acquisition (seeking advice when making important decisions), social bonding (reducing loneliness and reinforcing group bonds), and persuasion.

In building a theoretical framework for why we share political information, I draw on much of the that research. It might also seem sensible to use Berger's five functions model as a starting point for developing a theory of political information sharing. I do not do so, however, primarily because Berger's review is interested in word-of-mouth sharing of *product* rather than *political* information. Although there may be similarities with why people would share information about a product and why they would share political information, there are reasons to believe the motivations would only partly overlap. One shared motivation, for example, is to sculpt the impression others have of oneself; "impression management" is likely to play a role in the sharing of *any* information as it, indeed, plays a role in almost all social behavior. The incentives to share political information, however, may differ from sharing other information because, for example, political information is inherently contested and divisive; not only will people be divided on a piece of political information, but they are likely to be divided along group and political lines. Since most existing research on the sharing of information does not address contested information or group identification, it will lack what I believe to be key features of why users share political information.

A Proposed Framework for Understanding the Motivations of Political Information Sharing

As a baseline for understanding our information sharing behavior, I take as a premise that humans are goal-oriented. Ultimately, our goals are material; humans require resources to secure their survival and raise children that can then transfer their DNA into future generations. But, also like many species, our individual survival and reproductive success depends on the support of other humans. In pre-modern times to not have a secure web of social connections (not just family, but friends and groups) survival would be near impossible. Successful humans, therefore, must make the development and maintenance of social relationships a primary goal (Baumeister & Leary, 1995; Lieberman, 2013; Tajfel, 1974).

We can see that primacy of humans' need for social connection today not only in our behavior, but in our biology. The tremendous processing capacity of the human brain, indeed, evolved in large part to be able to negotiate complex human relations and develop language necessary to communicate with our fellow social beings (Baumeister & Masicampo, 2010; Dunbar, 1998). Although we can use that large neural capacity today to ponder non-social questions, our brain's default thinking mode today is still used to mull our social relationships (Lieberman, 2013). So important is the need to be connected to other humans, when we are excluded from social interaction our bodies register that isolation as physical pain (Eisenberger et al, 2003; Williams, 2007).

Group Impression

Impression Management and our Social Media "Press Secretary"

The social goals we have are varied – we may need to find a mate, develop a strong friendship or be accepted by the cool crowd – but our success in achieving any social goal hangs on how others perceive us. We are thus constantly in the game of - as sociologists say – managing our impression on others. To be a successful social animal we must mold and project a public "self" that has the characteristics appealing to current or prospective mates, friends or group members (Goffman, 1959; Kurzban, 2012; Leary et al, 2015).

The exact contours of a winning image will vary depending upon culture, age-group, domain (e.g. work or play), but generally any successful image has at least one of two qualities. For one, we will want to signal we are more “winner” than “loser”; others will want to connect with us if they see strength and ability, indicating that we will be more resource gain than drain. Second, it’s important that others trust us; a well-resourced friend is not worth much if they are either selfish or willing to betray us the moment it becomes advantageous (Neuberg & Cottrel, 2008). The ideal connection will be both a “winner” and a “good guy,” although humans are also generally willing to align with a powerful, resource-rich jerk or, conversely, a loyal loser. (No one, however, has time for a selfish chump.) For others to want to form connections with us, then, we will want to project an image that one is more gain than drain (i.e. competent, strong, rich, etc) and good (i.e. generous, fair and loyal).

Impression management colors and influences almost all our interactions, but it may be the dominant force in our engagement on social media (Nadkarni & Hofmann 2012). Compared to real life, social media sites permit a level of personal brand management that rivals Madison Avenue. In real life it is easy to slip and reveal our not-ready-for-prime-time selves, but on social media one can more carefully craft - and edit – the optimal public image (Hogan, 2010). That image can’t misrepresent who we are too drastically (Back et al, 2010), but it will be enhanced (Gosling et al, 2007) and massaged to fit a social ideal (Zhao et al, 2008).

It may seem that our friends craft that image consciously – with their obvious humble brags about their busy lives, filtered photos of glamorous vacation spots, and declarations of love for their “amazing” soulmates – but impression management is not always consciously driven. Much of the time that we project a favorable image of ourselves, it feels “authentic.” Recall that individuals need to perceive themselves in a positive light as both competent and morally good (Steele, 1988). Psychologist Robert Kurzban proposes the fact that we want to perceive ourselves as competent and good and that we also want to project an image of ourselves as competent and good is not a coincidence; theses parallel motivations are, indeed, two sides of the same image-conscious coin (Kurzban & Atkipis, 2007). More precisely, Kurzban argues that protecting and bolstering the personal self is not its own end. The self, rather, is a creation of our brain's "press secretary," a module that constructs an image and a narrative of ourselves that lets us successfully navigate our many social relations and achieve our social goals. The self we create is not for our own satisfaction; it is constructed to impress others. That we experience a self as a true entity and not as a mask constructed to win social games is a useful illusion; like any good press secretary, it does not need to know the self we project is a fabrication.[58] In that light, most of the time we don't need to do any additional work to project a socially-desirable image of ourselves that is distinct from our self-perception. Impression management thus will rarely feel like impression management – rather like an authentic expression of who we truly are.

Our Political Groups & the Imagined Audience

If we are trying to manage the impression others have of us on social media when we share political information, the next question is *who* are we trying to impress? We may also want to ask if there is a particular shade of "competent" and "good" we need to project. Sharing a political news story on Facebook is not like posting a photo of ourselves crossing the finish line of a marathon or sharing a campaign to raise money for suffering refugees; whereas the marathon post clearly projects "I'm a winner" and the fundraising campaign "I'm good," a post about global warming or a political candidate doesn't obviously project either image. Indeed, political posts seem impersonal and perhaps even selfless.[59]

[58] Tetlock and Manstead (1985) likewise proposed that impression management and "intrapsychic explanations" of behavior (including attitude polarization and other defensive mechanisms) may run in parallel and called for an integrative framework for both approaches to understanding behavior.

[59] If you ask people why they post news stories on social media they will themselves tell you that their actions are selfless and not at all connected to burnishing their personal image. As a preliminary test in my dissertation work (presented in full in Appendix II), I asked a number of social media users why they post political information; their most popular answer was that they wanted to keep their friends informed. But of course that answer is the one our "press secretary" would present – as it makes us look like good, selfless beings, just the kind of person you'd want to be friends with.

As a way of getting to the answer to those questions, we can look to two separate works of scholarship. First, I take the approach of information and technology theorists and ask if the affordances of social media can help us discern who users are communicating with. Next, social and political identity theory can also inform who we are talking to as well as give us insight into what impression we wish to make.

A prominent view among communication theorists is that technologies do not create new motivations and behavior, but rather merely provide a new medium to actuate motivations and behaviors that humans have always had (Gaver, 1991; Gibson, 1977). While not creating new behaviors, however, the features of a given communication technology - its affordances - will constrain or augment which motivations and behaviors are manifested. The telephone, for example, affords communication between two individuals which allows them to nurture social bonds, something that humans have always done but with the telephone can do at great distances (Bazarova & Choi, 2014; Boyd & Ellison, 2007; Ellison & Vitak, 2014).

Social media, however, is unusual in that it does appear to afford a new behavior - or rather it allows us to participate in a behavior that was traditionally limited. Before social media, few had a megaphone with which to communicate news to all their acquaintances. Perhaps we would pass along news to acquaintances in small groups at the dinner table or at meetings with co-workers, but the scope of information dissemination was limited. But with social media we all can all broadcast news, all the time.

Given that social media offers an activity previously limited to few, it is not obvious what - if any – behavior this new medium extends. It is also not clear to whom our communication is directed. When we talk to tens or hundreds on social media, what map in our brain are we using to understand what our goals and behavior should be? We know that individuals do not appear to be speaking to "all of one's friends and followers" when they post on social media (Acquisti & Gross, 2006). The category "everyone I know" does not have a place in our mental map. Instead, our brains must create a new map - or, more likely, kluge together a map from existing ones.

This puzzle – the difficulty of making sense of our own social media environment – has been identified by scholars who propose that social media creates a "contextual collapse" for users

(Marwick & Boyd, 2011). In real life we deal with situations within specific contexts - work meetings, family dinners, sports events - that define our goals and behaviors, and we adapt our behavior and image to the group we're with. This means that in real life, we get to be multiple people - or manage multiple identities. On social media, however, we don't have that option. Or, rather, we perplexingly don't use that option even though services like Facebook let us signal separate groups. Instead, all our contexts - our work, family, social groups - collapse into one.

We deal with this context collapse by constructing and speaking to an "imagined audience," one that is not necessarily reflective of reality (Marwick & Boyd, 2011). That imagined audience could instead reflect a group that exists in our real lives - e.g. our college friends, our professional colleagues - or could be imaginary, as when a poet might "conjure up a poetry fan community" while writing (Litt, 2012). Yet we know little about what makes up a user's imagined audience online, particularly when she posts political information.

I propose that we have *multiple* imagined audiences on social media, each loosely mapped onto real social groups in our lives. That is, our minds do not conceive of a new category of audience - the "everyone I know" audience; instead we rely on familiar social mappings in our brains. When we communicate on social media we do so as if we're speaking to one particular group - even though all of our groups can hear what we're saying.

Those social groups are more than just a set of individuals who have a shared trait or interest. Humans, again, are evolved to want to be part of groups and those groups play a central role in our psyche and behavior (Liebermann, 2013; Tajfel, 1974). It takes little for humans to identify with a group; merely telling someone they have a propensity to over-count dots, for example, can make individuals identify with and favor other "dot over-counters" (Tajfel, 1974). As members, we take on the beliefs, values and norms of the group (Suhay, 2015). The group becomes part of our identity and our brain adjusts to protect that identity (Tajfel & Turner, 1979). We see "our" group as distinct and superior to other groups and develop positive emotions toward fellow group members in contrast to feelings of animosity toward "others" (Hogg et al, 1995).

When it comes to politics, the imagined audience we conjure up when online could be any group for whom a political story is salient, including our racial, religious and cultural categories. We are also likely talking to our fellow partisans or ideologues. Much like our racial, religious and cultural groups, these political categories play a significant role in the identity of Americans. Political scientists have long recognized that American voters identify psychologically with their political party (Campbell et al, 1960). Green, Palmquist and Schickler (2004) argue that attachment is not to the party but rather to a "partisan group" for which, just as with a religious or cultural group, we create a "mental image" of the typical member that matches our own self-conception. We may hold a similar psychological attachment to an ideological identity (Conover & Feldman, 1981; Malka & Lelkes, 2010). Over the past few decades, those two political group attachments – to partisan and ideological groups - have taken increasing hold on our identities as they become not only more aligned with each other (liberals with Democrats and conservatives with Republicans) but also with our other social identities (Mason & Wronski, 2018).

Our Group Impression Goals & the Stories that Make the Right Impression

So what is the right image to project to those politically salient groups when we post political information on social media – and what kinds of news stories will do the trick? To answer those questions we first need to be more clear on what our goals are vis-à-vis our social groups.

I posit that the primary goals in creating an impression on our group are to a) secure our inclusion and b) to increase our status in that group. It is not enough that we simply identify with a group – the group must also accept us as a member. Again, the primary benefit of being part of a group is that it provides support – or even survival. We are therefore motivated to secures our acceptance in groups, but inclusion is not our only goal; it is also critical that we raise our status within that group, as status leads to not only more security but also to greater material benefits.

To be included in a group, first and foremost we will want to present as someone others can trust as a true and loyal member. Just as individuals will be wary of investing in a friendship with someone who will not reciprocate favors or, worse, betray one with little provocation, groups will be reticent to accept someone who may be a free-rider or, worse, infiltrator. To gain entrée into a group it is necessary to project clearly that you are a "true" member, not an imposter but rather the real, loyal, deal.

To gain status in one’s group, one will likewise want to reinforce an image of solid allegiance, but also project an image of individual strength and competence. In other words, to build alliances and supporters you will want to be someone others not only trust but also wish to ally with because they see such an alliance as leading to personal gain.

To demonstrate loyalty, we will want to post stories that signal we share the group’s values and beliefs or that display other signs of group allegiance, such as praising the ingroup or denigrating the outgroup. Users will thus be motivated to send those signals by sharing politically relevant stories (or images, quotes or memes) that:

● Affirm the group’s values, such as an inspiring quote or story.

● Either support the group’s beliefs or debunk the outgroup’s beliefs by, for example, providing new evidence that global warming is a looming threat or that illegal immigration is responsible for job losses.

● Burnish the character and stature of the ingroup or disparage the outgroup, as in a story about election polls that show the ingroup is winning or scandals that taint the opposition.

To signal strength and competence, a user has a number of options (Berger, 2014); she can present herself as one who has access to information or who has skills and capabilities potentially useful to the group. She could also demonstrate leadership or signal she already has high status. In the context of sharing political information on social media, the kinds of posts that could enhance our status include:

● Information useful to the group such as: breaking news about salient issues, political events that are relevant to the group, information about engagement opportunities, and facts or other evidence that are particularly potent in upholding group beliefs.

● Posts that indicate one’s knowledge, intelligence and cleverness, such as clever insights or quips about political phenomena.

● Posts that demonstrate leadership qualities, including calls to action or messages designed to inspire others.

Group Rallying

Responding to Individual and Group Threat; Calls of Distress and Moral Outrage

While the types of posts above will be familiar to any who have politically expressive friends on social media, the desire for inclusion and status may not account for all political posts. Two types of political stories commonly shared on social media that may be equally familiar may not be fully explained by group impression management: the “rant” and, its close cousin, the expression of moral outrage.

Communication theorists explain these types of posts as motivated by “emotional regulation”: that is, social media users feel so emotionally moved, they must express themselves to relieve those emotions (Berger, 2014). This explanation, however, while not inaccurate, only explains the immediate, emotional, cause of such behavior.

The reader may have noticed that I have neglected discussion of emotions as a motivation for sharing political information. This is not because I imagine emotions do not determine our behavior; to the contrary they play an essential role (Damasio, 1994). I, however, am more interested in the *goals* that motivate behavior. Emotions are not goals in themselves; rather they are the mechanism that both direct us to what “should be” our preeminent goal at any one moment and motivate us to act in the direction of that goal (Ellsworth & Scherer, 2003; Keltner & Gross, 1999; Marcus, 2000). To use the cliché example: when a tiger jumps out from behind a bush, we run away not to mitigate the fear we experience - rather, we run away so as not to become the tiger’s lunch. Likewise, when we experience hardship and feel distressed, we reach out to friends not to feel better but rather to get support. Yet, while mitigating our negative emotions is not our goal, emotions are the - absolutely essential - mechanisms that propel us to take action. If not for fear pushing us into a sprint, we *would* be the tiger’s lunch. In the theoretical framework I develop in this dissertation, however, my focus is on the ultimate, extrinsic goals that motivate sharing. I leave other scholars to study the emotional mechanisms that regulate and drive behavior in line with those goals (e.g. Brady et al, 2017; Stieglitz & Dang-Xuan, 2012).

Since emotion-based explanations do not fully tell us what motivates the rant or expression of moral outrage, we need to understand what purpose these strong emotions serve. I suggest that the rant and the expression of moral outrage serve distinct yet related goals. Those are to seek assurances of support when experiencing individual threat and to motivate group response to threats from outgroups.

The first motivation is straightforward. The reason we want to secure a spot in a group is precisely because at some point we will need the support or possible protection from that group. One would expect, given the affordances of social media, that when users experience threat, particularly in isolation, that they might use the platforms to send a “distress call” seeking support from their group. The rant can be thought of as such a distress call, one that articulates a case for why one is aggrieved (for example, by airlines that leave you on the tarmac or by fears of global warming) and elicits signals of support, reassurances that your community has your back or that they share your worries.

The second motivation - to coordinate group response to threat - is more speculative. The expression of moral outrage is widely observed as a central feature in the spread of incendiary news on social media (Haidt, 2016; Wayne, 2014). Researchers Brady et al (2017) likewise find empirical evidence that tweets which contain moral and emotional language get a retweet boost. One could explain the apparent popularity of “outrage” posts on social media as a form of group impression management: those who rage against the moral transgressions of the outgroup will, after all, make their group allegiance clear. And, yet, the power and seeming prevalence of the moral outrage post suggest that other motivating forces are at play.

Tooby and Cosmides (2010) provide a possible motivation. They suggest that expressions of moral outrage may play a unique role in group behavior; in times of threat from the outgroup, sharing moral outrage has the ability to motivate group members to take action against the outgroup, coordinating group behavior even without the benefit of a leader. Evolutionary theorists are often skeptical of such pro-social, seemingly altruistic behaviors that cost the individual for the good of the group. The authors argue, however, that the benefit of motivating others to act in concert would (in humans’ history) have had considerable benefits to individuals.

If the outgroup did, indeed, pose an existential threat to the ingroup, a cry of moral outrage that strengthens other members' resolve to fight the outgroup would be a relatively small price to pay.

Curating Ideologically Congruent and Extreme Stories

Above I have laid out a framework for understand what motivates and directs our choices in posting political information online. How can this framework explain users' propensity to share ideologically congruent and extreme news stories?

It is for the most part straight-forward how these motivations would lead to ideological curation. If the aim of sharing political information is to project an image of ourselves as loyal group members, to the extent that our ideological or partisan group is our "imagined audience" we will tend to share information that aligns with our – and our fellow ideologues' – beliefs about the world, policy and politics. That is, we will be ideological curators.

There are a few reasons these motivations might likewise compel users to post extreme news stories. Recall that "extreme" stories are those that tend to be dogmatic, emotionally- charged and tribal.

Not much needs to be said about why the motivation to stir moral outrage against the outgroup would compel users to share extreme information. To do so approaches a tautology. The expression of moral outrage requires both emotion (specifically anger) and an object of that anger (in this case, the political outgroup). We would expect then that when users want to express moral outrage against the outgroup, they will be more apt to share news articles that are anger-inducing and that strongly derogate the outgroup, two hallmarks of extreme news.

The motivation to signal one's allegiance to ideological and partisan groups might also drive social media users to share news stories that a) are unambiguous in their political message, b) promote ingroup favoritism, and c) are emotionally valenced. In sending signals of loyalty to the group, it is important that those signals are clear; a liberal user, for example, will prefer to share an article that unambiguously makes the case for gun control rather than a similar piece of information that is murky in its support. Likewise, sharing stories that portray ingroup members as virtuous or outgroup members as inferior (in competence or morality) sends a strong signal of group loyalty. Emotional content can also signal value as a group member; anger directed towards the outgroup, enthusiasm for ingroup wins or sadness at in ingroup loss all indicate group commitment. Thus, a motivation to signal group allegiance would make us more apt to share stories that are dogmatic, tribal and emotional – the three qualities of extreme political news.

In the next chapter I test to see if social media users do, indeed, have a preference for sharing stories with elements of extreme news – as well as other propositions made in this chapter. My goal here was to create a broad framework that might help us understand the motivations behind social media users' choices in posting and forwarding political information. This framework is a starting point that presents a plausible cohesive picture. The next step is to test its many propositions.

Chapter 7: Putting the Theory to the Test

In the previous chapter I presented a theory to explain social media users' choices in sharing political information. That theory made several claims. First I proposed that when sharing political information on social media, users are communicating with an imagined audience that reflects one of their political groups. I then posited two sets of goals that motivate their choices with that imagined audience: group impression management and group rallying.

As in offline life, users are motivated to project an image of themselves to their social groups that would make those groups want to have them as a member. The *group impression* they want to make is primarily that they are true, loyal members. Two ways to do that are to signal that one share's the group's beliefs and values and to cheer the ingroup while denigrating the outgroup. Users similarly will want their image to be of someone deserving high status in that group, which will likewise motivate them to share information favorable to the ingroup, but will also push them to indicate that they have higher competencies – knowledge, leadership qualities, etc.

Individuals aim to gain inclusion in groups for reason. Groups give an individual support and protection when threats are present; groups similarly are able to band together to cooperatively combat external threats to their collective wellbeing. When such threats arise, members are motivated to send out alarm calls, rallying others either to support oneself or to band together as a unified front against the threat. Thus when feeling threatened, either by policy-related fear or by the outgroup, *group rallying* will drive users to share content that beckons support and triggers a response against the outgroup.

Those twin sets of goals lead users to share news stories that are both politically aligned with the user and that are extreme in nature. To successfully signal group allegiance, at a minimum stories should be compatible with the values and beliefs of one's political groups; that is, stories should be politically congruent with the group's views. To be even more successful at signaling group allegiance, while in addition raising one's status in that group, those stories will also be dogmatic in upholding group beliefs, emotional in tone and direct in either extolling the ingroup or disparaging the outgroup, all hallmarks of extreme news. Likewise to rally group members in the face of threat, stories will use emotional language and accentuate tribal divides.

In this chapter I aim to find evidence to support several of those many propositions and predictions. Not having the time and space to test each facet of the theoretical framework, I focus on testing one set of goals – Group Impression Management – and one set of users' motivations - to secure inclusion in their political groups.

But even narrowing the focus to one motivation leaves several sub-elements and predictions to examine. In this chapter I test the following propositions:

- Users post political information on social media with an "imagined audience" that reflects one of the political groups they identify with;
- Users are motivated in part to shore up their inclusion in those groups:
- If the two claims above are true, we should expect that users will be inclined to share stories that are politically congruent and extreme. More specifically
    - Users will prefer to post stories that confirm their group's beliefs or make their group look good.
    - The more dogmatically a story confirms their group's view, the more a user will prefer to share that story.
    - Users will prefer to share stories that explicitly extol their group or denigrate the outgroup.

I take on those claims partly in reverse order, starting with the latter set of propositions - that users will tend to share stories that are politically congruent and extreme. Starting at the end makes sense for two reasons. For one, if we do not see that, indeed, users share information that is politically congruent and extreme in nature - and instead find that they are either neutral or they share stories that favor the outgroup - it would falsify the framework and we could end our hypothesis testing here.

Recall also that earlier we found that Twitter users do not seem to have a preference to share information that is favorable to their political groups; they appear rather to be unbiased in re-tweeting liberal vs conservative-leaning stories, which would suggest they are signaling neutrality. I suggested this may be the case because Twitter users had set up their information environment to reflect their views – obviating the need for them to selectively share stories sympathetic with their worldview. If, however, users were presented with equal shares of conservative- and liberal-leaning news stories, we might expect a bias in what they choose to share. I create such an environment in the set of studies presented below.

Selective Curation

Expectation/Proposition 1: **Users will prefer to post stories that confirm their group's beliefs or make their group look good.**

I begin with the prediction that users prefer to share stories that signal one shares the group's beliefs and is a reliable group cheerleader. In an online survey I present subjects with sets of five current news stories and ask them to place the stories in order of which they would be "most likely to share on social media." Among the stories presented, one is favorable to the left and one to the right – either because they uphold the group's beliefs and values or otherwise makes the group look good. If social media users who post political information are signaling their political group bona fides, when asked which story they would be most likely to share, we should see them place the story that is favorable to their party and ideology at the top of the list more often than the story favorable to the outgroup.

The News Stories Included in the Survey

In order to increase the realism of the study I present subjects with current news stories. By presenting real, current news stories, I remove the necessity for subjects to imagine what they might do at another time or with a fictional story - rather they can more easily and realistically imagine what stories they are most likely to share in the current political climate.

All stories came from Reuters, which I selected for two reasons. First, Reuters is a newswire service that does not have well-known ideological or partisan bias, so it avoids creating a possible confounder that subjects select stories solely based upon the political leaning of the source (Iyengar & Hahn, 2009). Using Reuters also minimized the sensationalism and emotional evocativeness of the stories. As much as could be hoped, then, subjects would select the stories based on their political content as opposed to source cues or other share-worthy qualities.

Stories were freshly collected each morning the study ran. I began each search by scanning Reuters' "U.S. News" page, looking for five types of stories: stories that favored Democrats; stories that favored Republicans; a nonpartisan hard news story (often an important international story, but also the occasional domestic natural disaster); a "fluff" story (celebrity or sports news); and one "random" story (any story that did not have an apparent partisan bias). I presented subjects an array of nonpolitical stories to choose from to both to create a more natural environment (a news consumer would be unlikely to only see partisan political stories when surfing the internet) and to avoid tipping subjects off to the fact that I was interested in their political choices. If I could not find four stories of each type in "U.S. News" I would scan Reuters' "U.S. Politics," "World News," "Technology," "Business" and "Entertainment" pages until I found a complete set. I aimed to find stories posted within the previous 24 hours, but occasionally it was necessary to go back two days to find twenty appropriate stories (and, rarely, three or more days).

For each article, I created an image that resembled what one would see on Reuters' web page, including the original headline provided by Reuters, the date, source (Reuters) and the blurb and photo Reuters used in listing their stories. In the cases where Reuters did not provide a blurb, I would use the first paragraph from their story. If there was no photo, I did a Google News search on the story topic and used the first photo the search provided.

In selecting stories that were favorable to the left or right, I aimed to pick stories that signaled group allegiance in one of two ways – either supporting the group's beliefs and values or championing the group in another way. Examples of the first type is story about Medicare spending staying within budget (confirming the Democratic view that public healthcare programs are affordable) and a story about record high incidence of heroin use (confirming a conservative view that crime levels are high). Stories that championed the ingroup include a report on Hillary Clinton leading Trump in the polls (seen in Figure 33) and an article about accusations of Clinton corruption. Of course, many stories cannot be exclusively placed into one category; a negative story about Obamacare can be both confirmation of conservatives' beliefs and a personal attack on a Democratic leader's legacy, while a story about a shooting massacre can likewise be supportive of gun control policy as well as an opportunity to malign the morality of the NRA.

Figure 33. *An Example Set of Stories in the Survey*

One set of stories subjects were exposed to in the study. Articles were presented in a similar way to how they would be seen on Reuters. Each set of stories included one article that was favorable to Republicans (in this example, the article about Obama's legacy) and one favorable to Democrats (Clinton's polling numbers). The three remaining articles were either a neutral hard news story (banking story), a fluff story (Game of Thrones) and a random story (Ohio State player).

I ran the survey in batches across five days with a new set of twenty stories each day. Using a large set of stories not only decreased the risk that I would select a poor batch of stories on one day, but also served to protect the overall results from the influence of any one story that was unusually “share-worthy” in some unintended way, for example by being particularly surprising or amusing.

Post Validating the Instruments

Because articles were selected the day of the survey there was not adequate time to pre-validate the target stories for their political valence. I instead depended on my own judgment in choosing stories the day of and post-validated the stories after the survey was run.

To do so I recruited Mechanical Turk workers (the same pool from which I drew participants for the survey) posting a 10 minute task for $1 to “Classify political and nonpolitical news stories.” Workers were asked to accept the task if they had “some familiarity with American politics and policies.” I excluded workers who had participated in the experiment.

Workers who accepted the task were first briefed in how to identify news stories that were “favorable to Democrats, favorable to Republicans or favorable to neither political group.” Coders were instructed a story could be favorable to one party or another for several reasons:

- “It makes the leaders or members of that party look good (eg. smart, capable, ethical, etc.)
- “It makes the leaders or members of the *other* party look bad (eg. stupid, incompetent, inethical, etc.)
- “It confirms the views of that political group, makes that group's priorities seem more important, or provides evidence to support the group's beliefs (eg. that global warming is a manmade problem (for Dems) or that immigration hurts American workers (for Republicans)).
- “It is otherwise good news for that party and/or bad news for the other party.”

Coders were also told that there might be other reasons a story could favor one party over the other and that they should use their judgment.

Workers then took a practice quiz asking them to label “A news story about gridlock, that blames Democrats and Republicans equally,” “A news story about a celebrity romance,” “A news story reporting a Democrat scandal,” and “A news story on new evidence that global warming is increasing.” Coders who mislabeled any of the items were shown the guidelines again and asked to retake the quiz. If a coder again mislabeled any item their classifications were not included in the analysis. Finally, at the end of the task coders were given a six-question political awareness quiz. I did not include classifications of coders who had more than two incorrect answers. Workers were each randomly assigned ten stories (out of the total 100 stories used in the experiment) to classify. After removing coders who failed either the training or political knowledge quiz, each of the partisan stories had between 4 and 13 ratings (with the average story receiving 8 ratings).

Looking at all coders who rated a story, more often than not there was agreement with my classifications, but not overwhelmingly so; average coder ratings agreed with me on the political valence of 25 of the 40 stories (so 63% agreement). However, much of the divergence lay in the fact that coders often found stories to be neutral, not favoring either party. If we restrict the analysis to only coders who said a story had a partisan slant, then average agreement between coders and the researcher rises to 89% of the time (or 33 out of 37 stories, with 3 stories all coders agreeing were politically neutral). This high level of agreement gives me confidence that if subjects in the experiment perceived a partisan lean to a story, it would be in the same direction as intended by the experimenter.[60]

Subjects, Design and Results

For this experiment, and all that follow in this dissertation, I recruited subjects from Amazon’s Mechanical Turk. Because Mechanical Turk workers are a convenience sample, it is not possible to say that the subjects in this experiment reflect the average social media user. In some ways, as discussed in Part I, Mechanical Turk workers are similar to typical social media users; for

[60] I similarly post-validated all the experiments presented in this chapter with similar results.

example they tend to be young and highly educated like the average Twitter user. In others they differ; they are more male than Facebook users and more white than Twitter users (Shearer & Gottfried, 2017). Mechanical Turk workers have, however, been found to be reliable subjects to the extent that they take surveys seriously, and partisan MTurkers (whom we are most interested in here) don't appear to be psychologically distinguishable from the average American partisan (Clifford et al, 2015; Levay et al, 2016).

I pre-qualified panelists by first inviting them to take a "1 minute survey" in which I asked their partisanship, ideology and social media use. For the survey experiment, I re-invited 1500 subjects who indicated they use social media, equal parts Republican, Democratic and Independent, to participate in a "10 minute survey." I ran the experiment using current stories on five separate days in June 2016, inviting 300 panelists each day. Over the five days, 513 subjects participated including 154 Republicans, 222 Democrats and 137 Independents.

In the survey experiment, subjects first answered a series of questions about their personality and social identity, followed by several filler questions about their news use and familiarity with famous figures. [61] Finally I presented them with four sets of five stories and asked "Which stories would you be most likely to share on social media? Please place them in order of 'most likely to share' to 'least likely to share.'"

In analyzing their answers, I looked at the frequency with which subjects said they would be most likely to share the story favorable to Democratic and Republican parties.

Partisans took the bait. Figure 34 shows the frequency with which Independents, Republicans and Democrats said they were most likely to share the Democrat or Republican friendly story. If subjects were to choose at random, they would place either story as a top pick 20% of the time. Both Democrat and Republican identifiers, however, place their party-favorable story as a top pick over 25% of the time. This is not only significantly greater than at random (in t tests, p

[61] One of those questions was about their party identity, which may have primed their partisanship. Yet I asked other identity questions so any such prime may have been dampened. In subsequent experiments party identity was not asked and effect sizes were even larger.

=0.003 and p=0.001 for Republicans and Democrats respectively), but is even more significantly distinct from the frequency with which they say they would share the story favorable to the other party (p<0.001 for both sets of partisans).[62]

Figure 34. *Selective Curation of Party Favorable News*

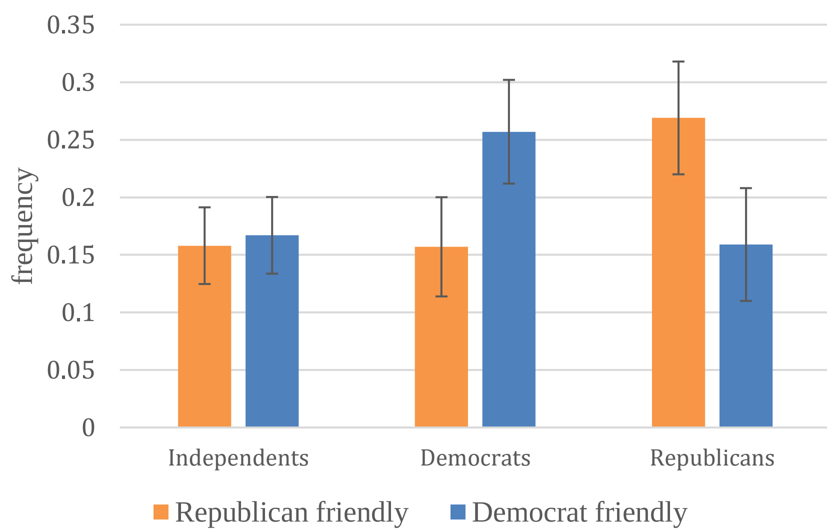

Frequency Independents, Democrats and Republicans selected the partisan friendly story as the one they'd be most likely to share on social media (out of five stories).

Subjects' aversion to sharing the counter-party story are also as we predict, but interestingly they share the counter-partisan story as infrequently as Independents do. This could mean one of two things. It could be that Independents share a distaste for partisan articles that equals the partisans' distaste for counter-partisan stories. The other possibility we cannot rule out, however, is that

[62] A question that may come to the reader's mind is if subjects say they would share these stories because they approve (or find the stories favorable) or because they dislike the stories. A policy victory for Republicans may, for example, be shared by a Democrat who wishes to rail against the policy. In this initial test I do not know how often – if at all – what appears to be cross-partisan sharing is, indeed, more in-party sharing. In later studies, however, I added a question at the end of survey experiments asking subjects if they chose the top stories because they "like" or "dislike" the news story. When I count disliked stories that are favorable to the out-party as in-party favorable stories, rates of selective curation rise even further.

these stories are un-shareworthy for other reasons; Independents and partisans alike may share them less because, for example, the stories are dull.

In all, the results above confirm the intuition that social media users are politically biased in the news they choose to share in the lab and lend support to Group Impression Management theory that users are motivated to signal allegiance to their political group.

A Note Online Surveys and External Validity

Although the design above allows us to create an environment where we can detect users' selection biases, it is not a natural environment and so presents limits into how much we can claim about subjects' behavior in real life. In other words, it is necessary to question the "external validity" of results we see in the lab.

Asking subjects in an artificial setting what they "would do" in the real world presents two types of external validity concerns. One is that subjects may be poor at putting themselves in an imagined situation and making predictions about their behavior. We would thus expect errors in their predictions. But while such errors may give us an imprecise view of users' true behavior, they do not necessarily bias our findings; with enough subjects those randomly distributed errors should cancel each other out.

The greater obstacle to inferring that in-lab results reflects out-of-lab behavior is the possibility that subjects share a *similar* bias in their predictions. In surveys that ask respondents if they are likely to vote, for example, people will both overestimate their commitment to voting and underestimate the obstacles that get in their way of making it to the polls on election day (Gollwitzer, 1999; Nickerson & Rogers, 2010). When they report to a survey taker that they are "likely to vote," they are really saying "I'd like to think of myself as someone who will vote." Similarly, subjects are well known to answer questions in "socially desirable" ways to look like upstanding members of society to the survey taker (Holbrook & Krosnick, 2009; Maccoby & Maccoby, 1954). Finally, if subjects can intuit the researcher's hypothesis they may try to behave accordingly to please the researcher, creating a "demand effect" (Orne, 1962).

The subjects in these surveys could be susceptible to such social desirability and demand influences, yet it is hard to say which direction those influences would sway them. In the case of social desirability, we might imagine participants like to think of themselves – and likewise to present themselves – as nonpartisan and unbiased (Klar & Krupnikov, 2016; Pronin et al, 2002). If that were the case then subjects may avoid saying they would share partisan stories, even though that is what they indeed do online. It is also possible that subjects guess the experimenter is left of center, which may incline them to share the more liberal story, regardless of their own political leanings. A third possibility is that subjects believe partisanship is socially desirable, in which case they might overstate their willingness to share partisan stories. Demand effects may also incline subjects to overshare the politically congruent story if they suspect that the study is about partisanship.

There are ways to mitigate social desirability and check for demand effects. A researcher could add directions that signal to subjects that any answer they give would be socially acceptable. It is also common to ask subjects at the end of an experiment what they thought the experiment was about to see if many subjects could intuit the experimenter's goals. Unfortunately this researcher did neither.

Nonetheless a social desirability effect presents little risk here. If subjects had an inclination to appear politically unbiased, they would suppress that bias, making the results above a conservative estimate of what happens outside the lab. We can also dismiss the possibility that subjects are trying to appear socially desirable to the – presumably – left-leaning experimenter. If that were the case then we would not see mirror effects for liberals and conservatives.

Subjects could, however, think it is socially preferable to appear partisan. If that were the case, then these results could be a reflection of social desirability rather than an indication of subjects' true sharing biases. Similarly it is possible that subjects have an in-lab tendency to share partisan-friendly information to please the researcher, if they sense that is what the researcher is aiming for. These effects would be a threat to the external validity of the results found above. Unfortunately, as mentioned above, I did not put in checks to minimize these threats so cannot dismiss them at this point. The next experiment I present, however, gives us reason to believe that subjects were not responding to demand effects, nor were they acting in socially desirable way to appear partisan. Stay tuned.

Selective Exposure?

While the results of the previous survey provide support for a Group Impression Management theory, they are also consistent with visceral selective exposure theory; subjects may simply be sharing stories they empathically know their friends would want to read. How, then, do we know that the observed behavior comes from users' motivation to impress their groups with their allegiance – rather than simply being a manifestation of visceral selective exposure?

Later in this chapter I test aspects of Group Impression Management theory that distinguish it from selective exposure (for example, by looking for evidence that users are communicating with an imagined audience). Before doing so, though, it is possible to see if selective exposure (and so visceral selective exposure) would produce the results we see above. In other words, we can check to see if selective exposure and selective curation are observationally equivalent; are our inclinations in what we choose to share the same as what we prefer to read ourselves? If they are, that would not necessarily mean that selective sharing and visceral selective exposure operate with the same underlying motivations; it is possible that two similar behaviors result from distinct motivations (one student might not cheat on an exam for fear of getting caught while another needs to uphold their moral self-image). But if selective exposure and selective curation are *not* observationally equivalent, then we have an indication that their underlying motivations may be distinct.

To see whether selective exposure and curation are observationally distinguishable, I recruited another set of participants from Mechanical Turk to take a version of the survey described above, this time with an experimental manipulation. All participants were again presented with sets of five stories, including one favorable to the left and one favorable to the right and asked to imagine that they saw those stories on social media. This time, however, I randomly assigned subjects to place the stories in order either of what they would be "most likely to share" or of what they would be "most likely to read" on social media. As before I ran the study over multiple days (three in this case), using current stories from Reuters. For this experiment I only invited subjects who had previously identified as either Republican or Democrat in a pre-qualification survey, since my question of interest was not how their behavior differed from Independents, but how it would shift depending on whether they were choosing which story to share or read. Over the three days 385 subjects completed the survey experiment, including 173 Democrats and 197 Republicans.[63]

Figure 35. *Comparing Selective Curation and Selective Exposure*

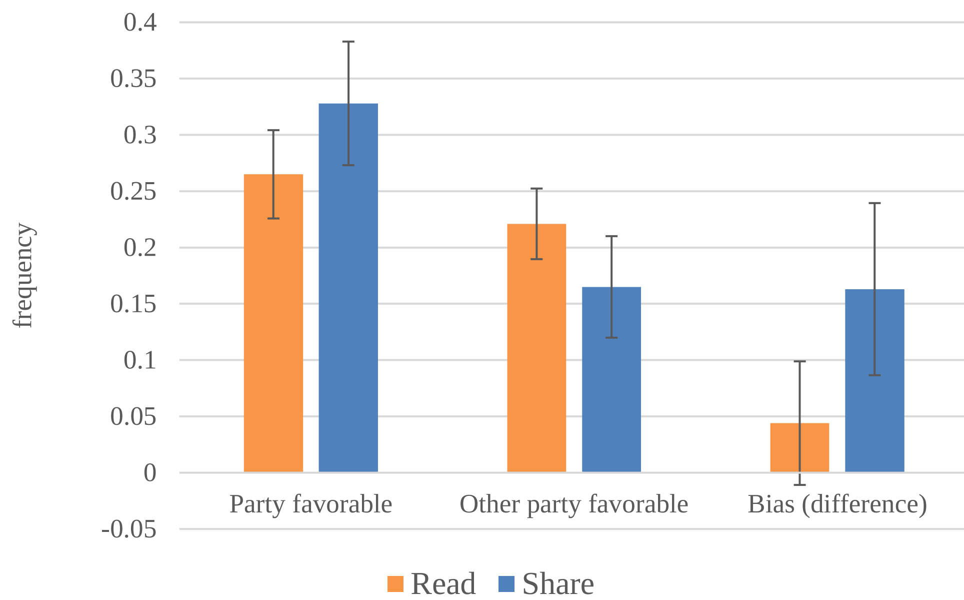

Frequency with which partisans selected their party-favorable story and the story favorable to the other party as the one they would most likely "read" or "share" on social media.

I aggregate the results, presented in Figure 35, for both Republican and Democrat subjects. As in the previous experiment, subjects showed a strong preference for sharing the news story that was favorable to their party while they are much less likely to share the story that is favorable to the other party ($p<0.001$). We likewise see a difference between subjects' preferences for *reading*

[63] 15 subjects were dropped because they could not be associated with their pre-qualification results.

party-favorable vs. other-party-favorable stories - 26.5% and 22.1% - yet that difference, shown in the third set of bars, only brushes against statistical significance (p=0.12). In contrast, the bias to *share* party-favorable stories over other-party-favorable stories is substantially stronger than the parallel bias to *read* congruent stories (p<0.001).

The results suggest that whatever motivations we have to selectively expose ourselves to political information, the motivations to selectively curate political information have a stronger footprint. In other words, while both behaviors may be fueled by the same motivations, when it comes to sharing, those motivations are stronger. But the results above suggest there may also be a qualitative difference. Recall that if subjects were choosing stories at random, they would select any story 20% of the time. Subjects in both the read and share conditions indicate a preference for the party-friendly story, although the preference for sharing is statistically greater than reading (p=0.01). When looking at the gap between how willing subjects are to read or share the *other*-party-favorable story, however, not only are subjects less willing to share that story (p=0.01) than to read it, subjects don't exhibit *any* aversion to reading that story; they are willing to read it at rates random chance would predict.[64]

This last result is in line with research in selective exposure, discussed above, which finds that although news consumers have a bias toward consuming congruent information, they do not actively avoid *in*congruent information (Garrett, 2009; Hart et al, 2009). Again, as motivated reasoning theorists posit, consumers are not only motivated to reaffirm their beliefs, but are also curious and motivated to have an accurate view of the world (Kunda, 1990). Even if they are not open to considering opposing views, they may be willing to read what "the opposition" has to say because they are confident in their ability to debunk their arguments. When we share

---

[64] There is a second way to read those numbers. It may be that, given subjects' strong preferences to share the party-favorable story, the chance that they share the "other party friendly" story is not 20%, but rather 17% (in the share condition) and 19% (in the read condition). With this angle, subjects in the share condition would be sharing the other party story at random, while the subjects in the read condition would be exhibiting an even stronger preference to select the other party story.

political information, however, even if we want to satisfy the curiosity of our friends, as Earl's "vicarious selective exposure" theory might propose (Earl et al, 2018), we may be less confident in our friends' ability to critically read information that threatens their beliefs, and so may be reluctant to share that information than we are to consume it ourselves.

But even with those nuances, vicarious selective exposure theory fails to fully explain the results. If selective exposure theory were able to tell us why we are willing to consume more disconfirming information than we share, it is not clear why we would be more eager to share than to read confirming information: if users are vicarious selective exposers, why would they be more desirous to confirm the beliefs of others than they would confirm their own beliefs? One way vicarious selective exposure theory might provide an explanation is if users thought their friends were all more partisan than they themselves were; this is possible and, yet, it would seem to stretch the power of empathy.

Group Impression Management theory better explains why it is users have a stronger tendency to share politically congruent information than they are to read it. If a user is motivated to present themselves as a loyal group member, they both have an incentive to avoid sharing information that may challenge that image in addition to share more information that burnishes that image.

It is easy to see why a social media user would not want to risk sharing information that makes the "other" group look good. While that user might, again, be willing or even curious to read an article that challenges their group's views, sharing such a story risks creating the impression that one is a doubter of the faith. One may even fear open backlash from other group members.

Our curiosity may drive us to consume not only challenging information but also non politically-valenced information; subjects in the experiment may thus be interested to read one of the fluff stories or one of the substantive stories that don't have a political message. In considering what to share, however, a fluffy story about a celebrity or a politically neutral story about a world event serve little purpose; they send no useful signal (other than that the user is shallow or boring - which may be other reasons they are avoided). The story that shows their party with a positive angle, however, is a sure hit; it will shed them in a light sure to gain the approval of their partisan friends.

As a final note, we can also use these results to set aside concerns about demand effects and social desirability in the earlier experiment. Recall that a potential bias in experimental research is that subjects will intuit the researcher's hypotheses and, because humans don't like to disappoint, respond in a fashion they believe will please the researcher. Subjects also might think it socially desirable to project an image of themselves as partisan. The results above make it unlikely such demand and social desirability effects exists for selective sharing. If one did exist, then we would have to assume that subjects intuited the researcher's hypothesis when asked about the stories they share, but less so about the stories they consume. Likewise they would only be acting in socially desirable ways (to the experimenter) when asked about sharing rather than reading. The pattern of results suggests subjects were responding to questions as they thought were true, not as they thought the researcher wanted to hear.

Testing Extreme Information

The experiments above support the proposition that users are inclined to share congruent information – that is, information that either confirms the group's beliefs or that makes one's group look good. It does not tell us, however, which attributes of a story make it shareworthy; is it the confirmation of the group's beliefs or is it, instead, the evidence that one's group is superior? The previous experiment also does not tell us if a story becomes more share-worthy if it more adamantly confirms the group's beliefs. In order to say if users have a preference to share extreme information, we would want to have answers to each of these questions. In the following experiments, I distill those extreme elements of a story – its dogmatic confirmation of a group's beliefs and its direct ingroup exultation - testing each in turn.[65]

---

[65] Again, in this work I do not test the third feature of extreme news – emotional valence – due to time constraints and the fact that other researchers have done work in that dimension.

Dogmatism

Once again, if social media users are motivated to show they share their political group's beliefs, they will need to share stories that align with those beliefs. But are all aligned stories equally effective at sending a signal of group allegiance, or will users – as I posit – have a preference to share stories that unambiguously confirm the group's values and beliefs over a story that is either lukewarm in its adherence to the group's views or that may leave room for ambiguity? A story that is nuanced or ambiguous in its viewpoint doesn't carry the necessary signal of group loyalty or - worse - risks projecting an image that one questions the group outlook, so we might expect social media users to share those stories less frequently than more dogmatic alternatives.

To test whether social media users do, indeed, share stories that are more clearly aligned with the ingroup's beliefs, we need some way of comparing relative levels of dogmatism and ambiguity in a news story while not inadvertently testing a confounding variable. In other words, we would want to see that when a news story dials up its unambiguousness, holding everything else constant, subjects with that group's identity are more likely to share that story.

To create such articles for comparison, I collected recent news stories on Google News using search terms for four policy issues for which the left and right are known to hold divergent views - gun control, climate change, immigration and minimum wage. For each issue I selected two stories that affirmed a liberal view (e.g. that global warming is a threat) and two that signaled a conservative view (e.g. immigrants are a threat). I then manipulated the headlines of those stories to either increase or decrease the unambiguousness of that signal. I did so in one of three ways: by altering either the certainty, salience or purity of the headline.

- *Certainty*: To make a headline more or less "certain" I simply introduced or removed words such as "maybe" or "certain." For example, I altered the CNN headline "Yes, climate change made Harvey and Irma worse" (certain in its connection between global warming and threat) to "Climate change may have made Harvey and Irma worse" (less clear evidence that global warming is a threat).

- *Salience*: While a news article may provide evidence to support a liberal or conservative point of view on an issue, the headline may or may not make the connection to that issue salient. For example, the Express headline “Earth could be plunged into mini Ice Age - and it could REVERSE global warming” makes no doubt that it is providing evidence that global warming is not a threat. In this case I simply removed the second phrase leaving “Earth could be plunged into mini Ice Age” to make the story less unambiguously about global warming.

- *Purity*: A news story can provide evidence that confirms an ideological point of view but at the same time provide countervailing information. The headline “Research suggests gun background checks work” clearly supports a liberal point of view about gun control, yet the actual headline (this time from NPR) “Research suggests gun background checks work, but they're not everything” sends a mixed message that is a less pure signal of support of the liberal line.

Using one of those three approaches, I created a total of sixteen pairs of stories (two for each political side for four issues), formatting them in the same way as the experiment above (headline, blurb and image) and pre-tested the degree to which each story confirm a liberal or conservative point of view. I recruited coders from Mechanical Turk and showed each ten of the 32 stories, and asked:

> Does this article confirm the beliefs of liberals or conservatives? (Does it provide information that supports one of the beliefs of liberals or conservatives?)
>
> - Strongly confirms liberals’ beliefs.
> - Somewhat confirms liberals’ beliefs.
> - Doesn’t confirm either liberals’ or conservatives’ beliefs.
> - Somewhat confirms conservatives’ beliefs.
> - Strongly confirms conservatives’ beliefs.

The more coders agree that a story confirmed a liberal or conservative view, the more “un-ambiguous” we could say it is. My aim was to identify eight pairs of stories (one liberal and one conservative for each of the four issues) that coders found statistically distinct. I also asked the coders to rate stories on several other dimensions - including whether they thought the story to be "important," "informative," or "interesting" - to make sure that any distinction within the pairs of stories was only related to how well it confirmed a left or right point of view. At least 23 coders rated each story.

The coders did not find all pairs distinct in ambiguity: I could only find four pairs of stories that were statistically significant (at a < 0.1 p-level using t tests). In order to have a larger set of stories to use as instruments, I included four additional stories that, although not statistically distinct, were distinct in the right direction.

Using those eight stories, I repeated the experimental design now familiar to the reader: subjects were recruited to look at sets of news stories and asked which they would be most likely to share on social media. I recruited subjects who had identified as Republican or Democratic in a pre-qualification survey, however I was interested in how their ideology (not partisanship) influenced their choices. In the experiment, the "treatment" was exposing ideologues to the less – as opposed to more - ambiguous story. I conducted the study in two batches; in each subjects only saw two of the four target stories. A total of 335 subjects completed the experiment, including 139 liberals and 160 conservatives.

Results

If ideologues are, indeed, apt to signal that they are true believers we would expect them share the less ambiguous story more often than the story that leaves room to question their adherence to an ideological view.

Looking at the frequency with which subjects say they are most likely to share the ideologically congruent story on social media in Figure 36, we can see that - across all eight stories - there was a marginally higher preference for the "strongly" unambiguous story than the "weaker" version, yet the difference is not statistically significant ($p=0.12$).

Figure 36. *Selective Curation of Weakly vs Strongly Dogmatic News*

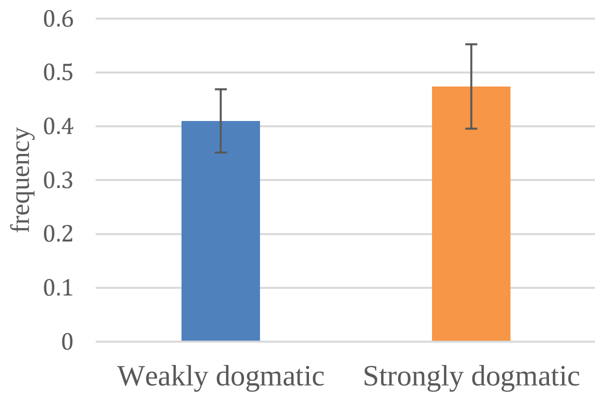

Frequency partisans share their party favorable story when its headline weakly or strongly confirms ingroup beliefs.

Moreover, if we break down the difference by individual story, is appears that much of the difference seen in Figure 36 can be attributed to one story; a story about immigrants that affirmed a liberal point of view. The other stories either show only a slim increase in the popularity when a story was presented in its "strong" version or, in two cases, saw a decrease in popularity. We thus should be doubly reticent to make anything of the overall finding.

Although the ambiguity hypothesis is not strongly supported by the experimental findings, a post-hoc analysis of the instruments used lends more credence that ambiguity may play a role in the popularity of stories subjects choose to share. In pre-validating the instruments, recall, I asked coders to say whether a story:

- Strongly confirms liberals' beliefs.
- Somewhat confirms liberals' beliefs.
- Doesn't confirm either liberals' or conservatives' beliefs.
- Somewhat confirms conservatives' beliefs.
- Strongly confirms conservatives' beliefs.

Using the coders' answers, I can assign each of the 16 stories a score from 1 (strongly confirms conservatives' beliefs) to 5 (strongly confirms liberals' beliefs), with a score of 3 indicating complete ambiguity. Reverse coding the scores of conservative stories, so both sets of stories are on the same scale from total ambiguity (3) to no ambiguity (5), and seeing which stories were most popular to share among subjects, we see in Figure 37 that there is a strong relationship between how clearly a story supports an ideological view and how frequently subjects sympathetic to that view say they would share it. Indeed, the clarity with which the stories support an ideological view predicts over half of the variance in their popularity.

Figure 37. *A Post-Hoc Analysis: Dogma and Selective Curation*

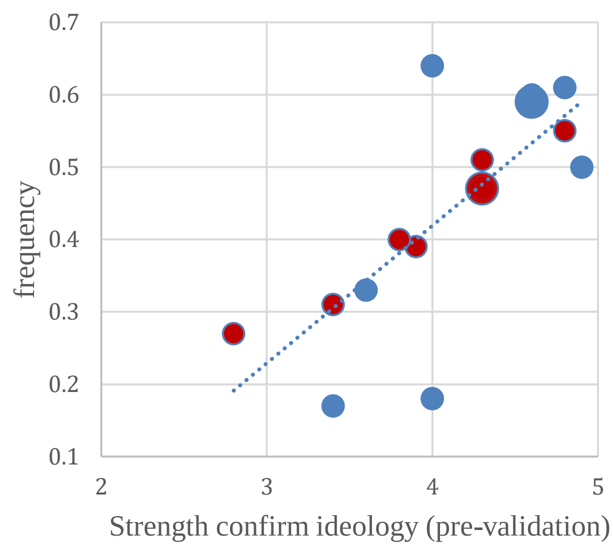

Frequency partisans are most likely to share their party favorable story relative to how clearly the headline confirms ingroup's beliefs, from 3 (total ambiguity) to 5 (minimal ambiguity). An observation below 3 indicates coders thought a story slightly confirmed the other party's viewpoint. The two larger circles indicate two observations on the same point. Republican favorable stories are red and Democrat favorable stories blue. R-squared = 0.56.

As a post-hoc analysis, we can only use this finding as exploratory research. Yet it points to a possible future experimental design. Rather than trying to construct identical stories that only diverge on levels of ambiguity, a future experiment could use multiple stories, independently scoring each for levels of ideological ambiguity and seeing which subjects say they are more likely to share.[66]

Extolling the Ingroup / Denigrating the Outgroup

A story that extols the ingroup or denigrates the outgroup sends a clear signal of group solidarity - and thus we might expect it is more appealing to share than a similar story that is neutral towards one's in and out groups.

To test if this is the case, I use an experimental design like the one above again manipulating the headlines of real news stories, but this time altering articles to clearly praise the ingroup, tarnish the outgroup or, as a control, be group neutral. I looked for stories that extolled the ingroup or denigrated the outgroup in different ways, either by distinguishing a group's inherent traits, applauding an achievement of the ingroup or noting an immoral act of the outgroup:

- *Superior group traits*. One way to extol the ingroup is to share stories that report the group has superior traits. Luckily there is no shortage of these types of stories circulating online; I easily found articles that reported on differences between liberals and conservatives on their levels of intelligence, attractiveness, bigotry, anger, charitability, happiness and gullibility. I chose stories about two traits: one story linking low I.Q. to conservatism and one reporting on the superior attractiveness of conservatives. I then manipulated the stories' headlines and blurbs so that they attributed the inferiority/superiority to either group or to neither. I chose intelligence and attractiveness because both seemed plausibly attributable to either group and because they varied in their political salience and importance. Intelligence is a trait that might be considered an important and relevant quality for praising a political group; one could argue, then, that evidence that affirms a group's intelligence is relevant political information. A group's level of attractiveness, in contrast, should be politically irrelevant; thus if we see subjects

[66] Although such a design would not completely rule out that other factors influenced a story's popularity.

share stories that tout their group's good looks, this would indicate unadulterated ingroup cheerleading.

Figure 38. *Example Stories: Superior Group Traits*

Liberals really are better looking research says

Reuters | Mar 18

A recently published study in the Journal of Public Economics concludes that the attractiveness of a candidate correlates with their politics. They find that politicians on the left are more good looking in Europe, the United States and Australia.

Conservatives really are better looking research says

Reuters | Mar 18

A recently published study in the Journal of Public Economics concludes that the attractiveness of a candidate correlates with their politics. They find that politicians on the right are more good looking in Europe, the United States and Australia.

Conservatives and liberals equally good looking research says

Reuters | Mar 18

A recently published study in the Journal of Public Economics concludes that the attractiveness of a candidate does not correlate with their politics. They find that politicians on left and right are equally good looking in Europe, the United States and Australia.

- *Credit for positive outcomes*. Another way to cheer the ingroup is to share information that applauds the group's positive achievements; if a political group is shown, for example, to improve the economy or reduce crime, it would affirm that the group has superior policies and leadership. They also show that the group is "winning." Ideally, in an experimental manipulation I could take a similar approach to the "trait" stories above; that is, use the same stories for both political groups and merely alter which group could be credited with a policy achievement. The challenge was finding stories that could be equally plausible regardless of attribution; since subjects would be likely to know who is in political power and the issues that parties care about, it would not be believable to read, for example, “Democrats improve economy”[67] or “Republicans reduce global warming.” To find articles that could be plausibly manipulated, I looked for stories that a) were about policies that are *not* known to be owned by one party and b) did not require one party to be in power for them to claim credit. What fell into those categories were two stories about bipartisan efforts to introduce legislation on cybersecurity (shown) and battling opiate addiction. Another advantage of using these stories is that I did not have to present subjects false information; subjects merely saw partial information in the manipulation.

Figure 39. *Example Stories: In-party Credit*

Democratic lawmaker makes new push on cybersecurity

Reuters | Mar 19

Sen. Kamala Harris (D-Calif.) said Monday that she is planning to introduce an amendment to a bill reauthorizing the Department of Homeland Security (DHS) that would help states modernize their election systems.

Republican lawmaker makes new push on cybersecurity

Reuters | Mar 19

Sen. James Lankford (R-Okla.) said Monday that he is planning to introduce an amendment to a bill reauthorizing the Department of Homeland Security (DHS) that would help states modernize their election systems.

Lawmakers makes new push on cybersecurity

Reuters | Mar 19

Sen. Kamala Harris (D-Calif.), Sen. James Lankford (R-Okla.) said Monday that they are planning to introduce an amendment to a bill reauthorizing the Department of Homeland Security (DHS) that would help states modernize their election systems.

- *Immoral behavior*. Finally, I used stories that impugn the morality of the outgroup. These came in two flavors: political scandals (both sex and corruption) and stories that attributed racist actions to members of the outgroup. As with the “credit giving” stories above, there were challenges to finding suitable current news items that I could easily

[67] At the time of this study, Republicans held Congress and the presidency.

adjust for both Democrats and Republicans. For example, a story about a Republican mayor calling immigrants “raccoons” may have been implausible if I insisted he was a Democrat. I was also concerned about falsely attributing negative behavior to political groups; even though subjects would be debriefed at the end of study, false beliefs are known to linger even after corrected (Nyhan & Reifler, 2010). Instead of using the same story for each political group, I selected different stories for the liberals/Democrats and conservative/Republican subjects and manipulated the headline to merely note the party membership of the immoral actor or leave their party information out. It was also necessary to choose politicians and news items that were significantly obscure so that subjects would be unlikely to know their political affiliation. This choice necessitated using local politicians with little power, which risked weakening the intervention.

Figure 40. *Example Stories: Outgroup Immorality*

Republican lawmaker allegedly used taxpayer money for sex with hooker

Reuters | Mar 6

A Utah lawmaker who voted for tougher penalties for prostitution has resigned amid allegations that he used taxpayer dough to pay for hotel rooms to hook up with an online escort, according to reports.

State lawmaker allegedly used taxpayer money for sex with hooker

Reuters | Mar 6

A Utah lawmaker who voted for tougher penalties for prostitution has resigned amid allegations that he used taxpayer dough to pay for hotel rooms to hook up with an online escort, according to reports.

Nashville Mayor resigns after affair, pleading guilty to theft

Reuters | Mar 6

Nashville Mayor Megan Barry resigned on Tuesday after pleading guilty in state court to theft in connection with an extramarital affair she admitted to having with the head of her security detail.

Democratic Mayor resigns after affair, pleading guilty to theft

Reuters | Mar 6

Nashville Mayor Megan Barry resigned on Tuesday after pleading guilty in state court to theft in connection with an extramarital affair she admitted to having with the head of her security detail.

Results

In addition to recruiting subjects with known partisan identities, I also included Independents; I wanted to be sure that any increased preference for sharing the ingroup/outgroup story was specific to partisans and not a general preference for stories that specify party or ideological affiliation. As with the studies above, I recruited subjects from my Mturk pool, offering $1 to take a "10 minute survey." In total, 390 subjects took the online experiment, including 132 Republicans, 130 Democrats and 128 Independents. Each subject saw five sets of five stories; in three of those sets they saw one of the target stories, either a group neutral story or a story that extolled or denigrated one group.

Overall, subjects with party affiliations showed a strong preference for stories that puffed up their party or associated ideological group (p=0.01) compared to neutral stories. This preference was not shared by independents; in fact, independents showed an antipathy toward stories that differentiated ideological groups). A difference-in-difference analysis between independents and partisans shows a 12 percentage point difference in treatment effect (p<0.01).

Figure 41. *Extolling the Ingroup / Denigrating the Outgroup*

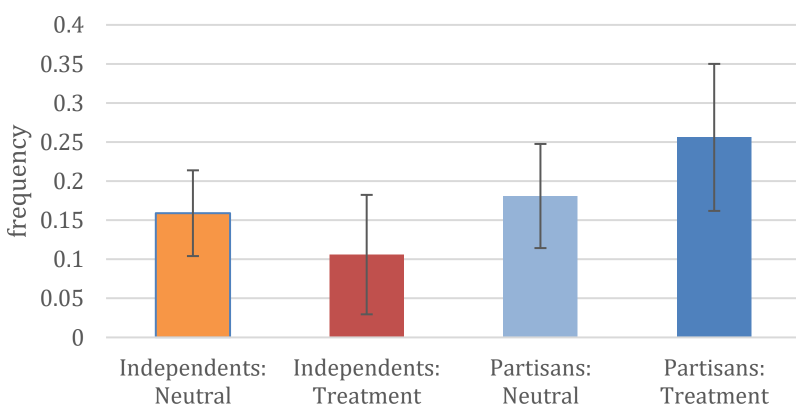

Frequency that partisans and independents share stories that (for partisans) extol the ingroup or denigrate the outgroup (treatment) or are party neutral.

As with the previous experiment, not all stories had similar treatment effects. The overall effect seen above appears to be driven largely by the stories that differentiate between the inherent traits of members of ideological groups.

Figure 42. *Extolling the Ingroup: Superior Traits*

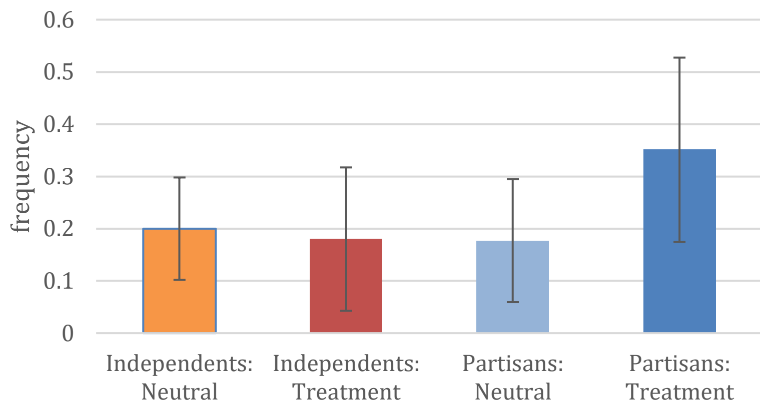

Frequency that ideologues and independents share stories that (for ideologues) report on the superiority of their ingroup (treatment) or are group neutral.

When liberals and conservatives are reported to have similar levels of intelligence or attractiveness, partisans and independents alike show no preference toward sharing those stories. If a subject's ingroup, however, is reported to have higher intelligence or attractiveness, that subject will be substantially more likely to say they would share that story ($p<0.01$). This result cleanly confirms that ideologues have a preference toward sharing stories that extol their ingroup.

Figure 43. *Extolling the Ingroup: Ingroup Credit*

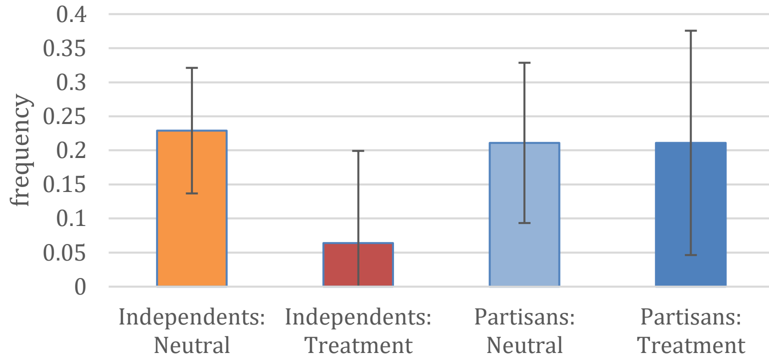

Frequency that partisans and independents share stories that (for partisans) give their party credit for policy gains (treatment) or are group neutral.

The same cannot be said for news stories that credit the ingroup with legislative action. Partisans are equally likely to say they will share a story that reports on a bipartisan legislative effort as when that effort is credited solely to their party. They are, however, significantly more likely to share a story that credits their party than an independent would (p=0.01). One way to explain these unexpected findings is that partisans and independents alike are fairly neutral about these stories, since they share them at a rate equivalent to random, with the exception for independents who have an *antipathy* toward sharing stories that credit one party over the other. Another possible interpretation is that a story about one lawmaker or one party introducing a bill is inherently dull, but such a story sends a share-worthy signal for a partisan when the subject's party member is responsible. In that light, these results could confirm our expectations that partisans have a predilection toward sharing stories that signal their group's achievements.

What remains a puzzle is why both partisans and independents are likely to share a story about bipartisanship. One possible answer is that independents and some partisans share a desire to appear balanced and nonpartisan (Klar & Krupnikov, 2016).

Figure 44. *Denigrating the Outgroup: Morality*

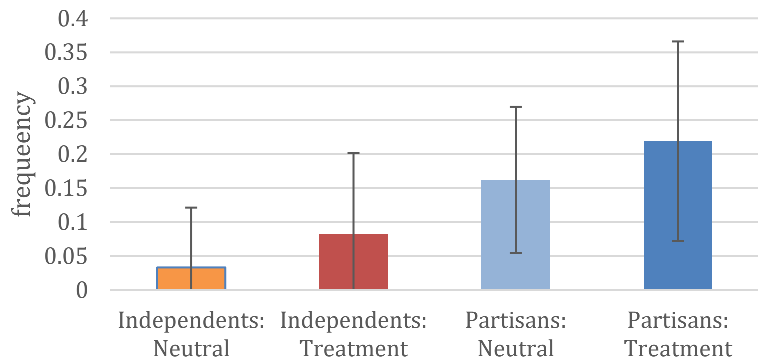

Frequency that partisans and independents share stories that (for partisans) impugn the morality of the outgroup (treatment) or are group neutral.

Finally, subjects' preferences for sharing stories that point out the moral failings of the outgroup are only somewhat in line with our expectations. Partisans do show a preference for sharing stories about corruption and bigotry when it is clear that the perpetrator is an outgroup member, yet that difference does not reach statistical significance (p=0.23). What is unexpected is Independents' comparative aversion to sharing any morality story (p=0.02, just looking at the group neutral story).

Once again, there are at least two possible interpretations for these results. Partisans may feel indifferent about the signal these stories send (again, they share them a rate near random), while Independents have an aversion to spreading negative political stories and signaling that they care about mucky politics (Hibbing & Theiss-Morse, 2002).

Another possible explanation for the results in Figure 44 is that, like the stories about bill introductions above, these stories are inherently dull unless a user is a partisan and there is an opportunity to impugn the morality of one's outgroup. That is easily done when the party of the moral transgressor is explicitly noted. It may also be possible, however, if the party of the transgressor is easily intuited to be that of the outgroup. Looking at the stories I used for moral transgressions, that could be the case for the Republican stories I used (as I notice in hindsight). As the reader can see in the instrument used for a Republican lawmaker in Figure 40, even without being told he was Republican subjects knew that he was a state legislator in Utah, which is a strong clue of conservative and Republican leaning. The second instrument I used for Republican moral transgressions may have been equally obviously Republican (Figure 45); in this case, subjects saw an image of a white man who called undocumented immigrants "raccoons," which may have elicited an association with conservatives.

Figure 45. *Group Neutral Story for Republican Immorality / Racism*

Mayor resigns amid backlash over post comparing immigrants to raccoons

Reuters | Mar 18

Mendham Township Mayor Rick Blood resigned after igniting a firestorm with his Facebook post comparing undocumented immigrants to raccoons in a basement.

Group neutral instrument that may not have been so group neutral (i.e. subjects may have guessed the Mayor in question was conservative and/or Republican).

The instruments used for immoral Democrats, in contrast, would not have been as obviously Democratic if their party had not been explicit. The mayor of Nashville in Figure 40 is a woman and city mayor, both associated with liberalism, but she is also a Southerner, which might send mixed signal about her party affiliation. The second instrument used, about Democratic senators supporting a possibly racist bill (Figure 46) would likewise not be associated with Democrats if they were not explicitly noted.

Figure 46. Group Neutral Story for Democrat Immorality / Racism

Senators Back a Bank Bill That Could Hurt Black Homebuyers

Reuters | Mar 19

A provision in the Economic Growth, Regulatory Relief, and Consumer Protection Act would exempt the large majority of mortgage lenders from key disclosure requirements that help the government identify racial discrimination and enforce fair housing laws.

Group neutral instrument that was more likely to be truly group neutral (i.e. subjects would be less likely to guess senators were liberal and / or Democrat).

It may be the case, then, that while conservative subjects would need an explicit cue to identify Democrats as the moral transgressors in the articles, liberal subjects would need no such cue. If we compare the sharing preferences of Democrat and Republican subjects in Figure 47 we see this could be the case. Subjects who identify as Democrat were nearly equally willing to share the story about an immoral Republican regardless of whether his party was explicitly identified or not. Republican subjects, however, were substantially more likely to say they'd share the morality story only once they were told the perpetrator was a Democrat (p=0.04).

Figure 47. *Morality - Comparing Democrats & Republicans*

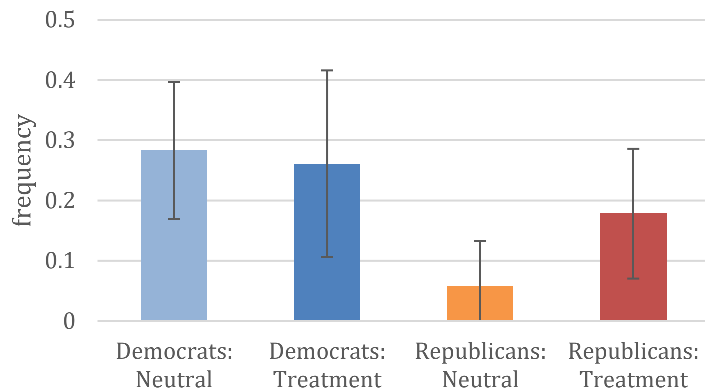

Frequency that Democrats and Republicans share stories that explicitly impugn the morality of the outgroup (treatment) or are group neutral.

In sum, we find evidence that partisans and ideologues are more inclined to share stories that exalt one's ingroup or conversely disparage the outgroup. This is most cleanly and clearly seen in stories that attest to the ingroup's inherent superiority; stories that affirm the ingroup is either more intelligent or good-looking than the outgroup are not only more likely to be shared by ideologues than stories that say groups have parity, but they are more likely to share them than ideologically moderate subjects. Partisan subjects are also more likely than independents to share a story that gives their party credit for introducing legislation. Finally, partisans are more likely than independents to say they'd share a story about a morally corrupt politician from their outgroup, even though (possibly given this researcher's poor choice of instruments), we only see such an effect with Republican subjects.

These experiments both give us a clearer picture of what drives a social media user to share politically congruent information as well as why they may be more apt to share extreme stories. The ideologues in these studies are unabashed tribal cheerleaders. We can see that most clearly in the stories about groups' inherent traits; if a story explicitly notes the superiority of their ingroup – even regarding a quality that should be irrelevant to one's political worldview – subjects are happy to put it on their social media wall. Similarly, though to a lesser degree, if a story credits the ingroup or impugns the morality of a member of the outgroup, it is seen as more shareworthy. We also saw evidence, though weak, that stories will be more desirable to share the more they dogmatically confirm the ingroup's beliefs.

Expectation 2: Social media users are communicating with an "imagined audience" that reflects one of their political groups

The experiments above confirm the intuition that social media users prefer to share stories that are favorable to their political groups. They likewise confirm that stories are more share-worthy if they exalt the ingroup or denigrate the outgroup or, to a lesser degree, if a story is more dogmatic in confirming the ingroup's beliefs. We also saw that users' sharing choices were not a mere mirror of "selective exposure." These results are compatible with the Group Impression Management theory I have presented – i.e. that social media users are motivated to project an image of themselves as true and loyal group members. This is a good start. If users had failed to exhibit a tendency to share information favorable to their group or if such a bias could be explained by the same forces that produce "selective exposure," it would be hard to claim that users were motivated to signal to group members that they are worthy members. But we still lack evidence that users are, indeed, signaling to their group members and that they do so to bolster their group inclusion.

In this next study I test the claim that users are signaling - or communicating - with their political groups as their "imagined audience." To do so, I lean on theory and methods not from the scholars who study imagined audiences, but from researchers who study social identity. As discussed previously, social identity theory posits that part of our individual identity is formed by the social groups that we belong to. Those groups can be multiple and varied, and can include racial, religious, professional, and cultural groups (Huddy, 2001; Tajfel, 1978).

Those social identities can also wax and wane in prominence depending on our current environment (Cameron, 2004; Hogg & Reid, 2006). Researchers take advantage of this multiplicity and lability of social identity to study its influence; by priming one identity over the other in an experimental setting it is possible to see, for example, how social identities influence political preferences (Klar, 2013; Jackson, 2011).

I borrow those interventions to see if I can prime one of a subject's social groups and, in so doing, shift their "imagined audience" and alter the stories they choose to share. In using this manipulation, I am assuming that by raising the salience of a social identity, I likewise am heightening the presence of one social group in a subject's mind. This is a bit of a leap, but not an insensible one. When, for example, I think of myself as a "New Yorker," I can only do so by thinking of myself as a part of the group "New Yorkers" (Turner, 1975). One's group identity would thus naturally elicit consciousness of one's group. My expectation is that if subjects are primed to have one of their social groups more present in their mind they would be more apt to think of that group when selecting a story to share.

In two separate studies I recruit subjects who have dual social identities to test that hypothesis. One set of subjects identify both with a non-white minority group (Black or Latino) and as Democrats. Another set, in which I include only white subjects, identify both as partisans (either Democrat or Republican) and, given their location in the US, I expect will also identify as American. (And, indeed, 95% of subjects in the experiment say they identify as American at least "a little.")

In the case of the experiment with white subjects, I use the identical instrument as in previous experiments; subjects once again are presented with sets of stories that each include one article that is favorable to Democrats and one favorable to Republicans and are asked which story they would be most likely to share. Subjects are randomly assigned into a control group or one of two treatment conditions in which they are primed either with their partisan identity or their

American identity. My expectation is that when subjects' partisan identities are primed, their partisan group will become more salient and they will be more likely to share the story that reflects well on their party. When their American identity is primed, in contrast, I expect that their partisan inclinations will be mitigated; with the super-ordinate social group in mind - one that encompasses Americans beyond just their partisan sympathizers - they will be less inclined to signal their identity as partisans and might even share stories that indicate they are open to more diverse views. (For a similar approach investigating selective exposure see Levundusky, 2018.)

For the subjects who identify with one of the minority groups I use a somewhat modified instrument. Subjects likewise saw sets of five news stories among which one is favorable to Democrats, but instead of seeing a Republican-friendly story subjects see an article that is relevant to their minority identity. For the Democratic stories I use the same type of stories as in previous studies; that is, recent stories from Reuters that either confirm a Democratic point of view or make Democrats look good. For the articles intended to be relevant to the subjects' minority identity, I selected and pre-tested three types of news stories: one that showed the minority group in a strongly positive light (e.g. a story about an African American Olympian), one that was threatening to the group (e.g. a story about low test scores among Latino students) and one that was relatively neutral (e.g. a story about minorities and the housing market). Similar to the experiment with white subjects, the expectation is that the story subjects share will vary depending upon whether the subject's Democratic or minority identity is primed.

Finally, for both studies I include a parallel set of conditions in which I ask subjects which stories they would be most likely to *read*. Like the earlier experiment which included a "share" and "read" condition, I want to be sure that any effects we see are related to subjects' motivations for *sharing* information. If we were to observe that subjects are more apt to share party relevant stories when primed with their partisan identity, for example, it might not be the case we have shifted their "imagined audience," but rather simply that their party becomes more salient, drawing them more to those stories. By including a read condition we can test – and possibly discount - that alternative hypothesis.

Results: Priming Race and Party Identities

In the experiment with minority subjects, I initially primed subjects' identities simply by asking if they identified with a minority group (racial prime) or party (partisan prime). I recruited Mechanical Turk workers who had previously identified themselves as Black or Latino in an earlier qualification survey; 67 Black and 60 Latino subjects completed the survey experiment.

I present the results of the Black and Latino subjects in Figure 48.[68] What is first worth remarking is the overall popularity of the minority-relevant stories; subjects say they would share those stories 42% of the time in the control condition, more than twice the rate one would expect if they were choosing randomly. Stories that are pro-Democratic are relatively unpopular in the control condition, with subjects selecting those stories about 15% of the time. When primed with their partisan identity, however, subjects increase their preference for sharing the Democratic stories, though the difference does not quite reach statistical significance (p=0.09). Such a result suggests that subjects' "imagined audience" has shifted; they still share stories that are relevant to their minority social group, but with less frequency.

[68] Asian subjects' results were flat; a possible explanation is that, as Chong (YEAR) shows, pan-Asian identity is weak.

Figure 48*. Effects of Priming Identity on Shifting the "Imagined Audience" I*

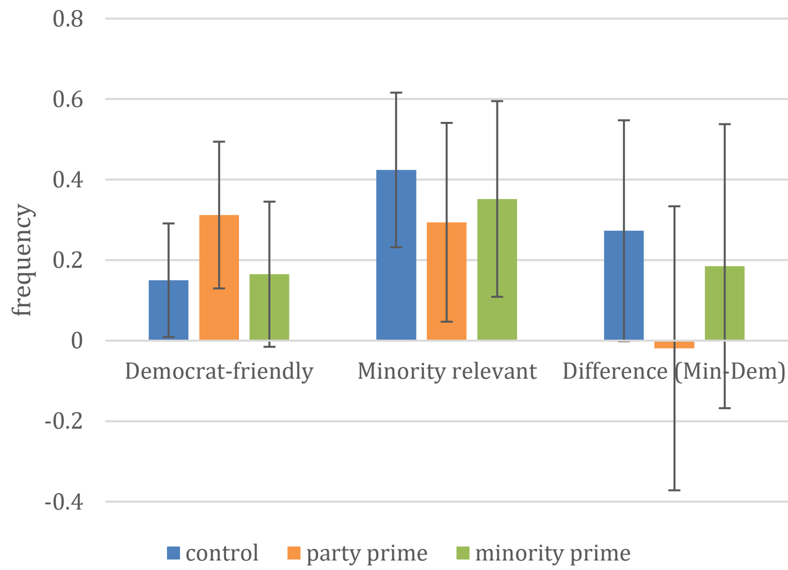

Frequency minority subjects (all of whom are Democrats) selected the story that was favorable to Democrats ("Democrat-friendly") or salient to their minority group ("Minority relevant") when they were primed with their party identity ("party prime"), their minority group identity ("minority prime"), or with no prime (control). (N=64)

What we do not see, however, is a shift in the stories they share when primed with their minority social group; if anything we see a decrease in the frequency they say they'd share the minority relevant story. But the absence of an increase may not be too surprising; the popularity of the minority stories in the control condition suggest a "ceiling effect" could be at play. For racial minorities, race is a dominant group identity (McClain et al, 2009); priming racial identity will thus have little effect in increasing the relevance of that group. In other words it may not be possible to make subjects' racial group much more salient than it is in the control condition.

While the results above are suggestive, the sample size is so small that we cannot confidently say that there is a shift in subjects' "imagined audience." To see if the results are a fluke, I re-run the experiment, this time recruiting Mechanical Turk subjects through Turkprime, a service that lets researchers recruit subjects with specific demographic characteristics. In this second experiment I also strengthen the party and race primes by adding a secondary question asking subjects how important that identity is to them.

Combining the results from that experiment with the first experiment (N=195) we see similar results. Again, when subjects with dual Democrat and minority identities were primed with their party identity they were more apt to share the story favorable to Democrats, and the difference approached statistical significance (p=0.06) than in the first study. Likewise, respondents were yet again not more likely to share the story salient to their minority group when primed with that identity (p=0.88).[69]

Figure 49. *Effects of Priming Identity on Shifting the "Imagined Audience" II*

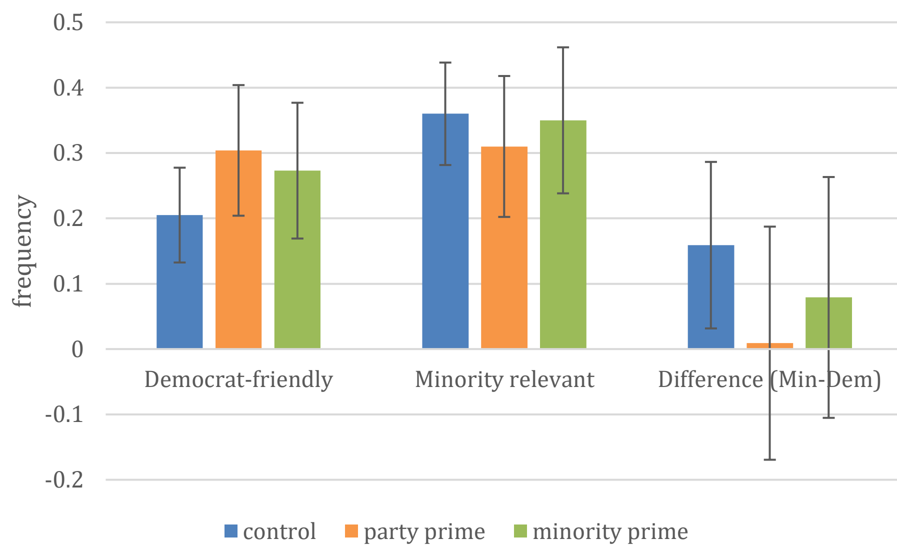

Frequency minority subjects (all of whom are Democrats) selected the story that was favorable to Democrats ("Democrat-friendly") or salient to their minority group ("Minority relevant") when they were primed with their party identity ("party prime"), their minority group identity ("minority prime"), or with no prime (control). (N=260)

The results above suggest that when primed with their partisan identity, subjects are more apt to share party relevant stories, but does that imply a shift in their "imagined audience"? To answer that question we can compare the priming effect on what subjects say they will share to what they say they would *read*. Looking at such a priming effect in Figure 50 we see no greater

[69] The analysis of this second study varies slightly from the first in how subjects' partisanship is identified. Because I am not able to pre-select Turkprime subjects on their partisanship (but only did so on their race) it was necessary to identify their partisan leaning after they completed the study.

willingness to read the Democratic story when subjects are primed with their partisan identity. A statistical analysis (looking at the interaction effect of "party prime" * "share vs. read") confirms that the effect of priming partisan identity is unique to the share conditions (p=0.04).

Figure 50. *Effects of Priming Identity on Selective Exposure*

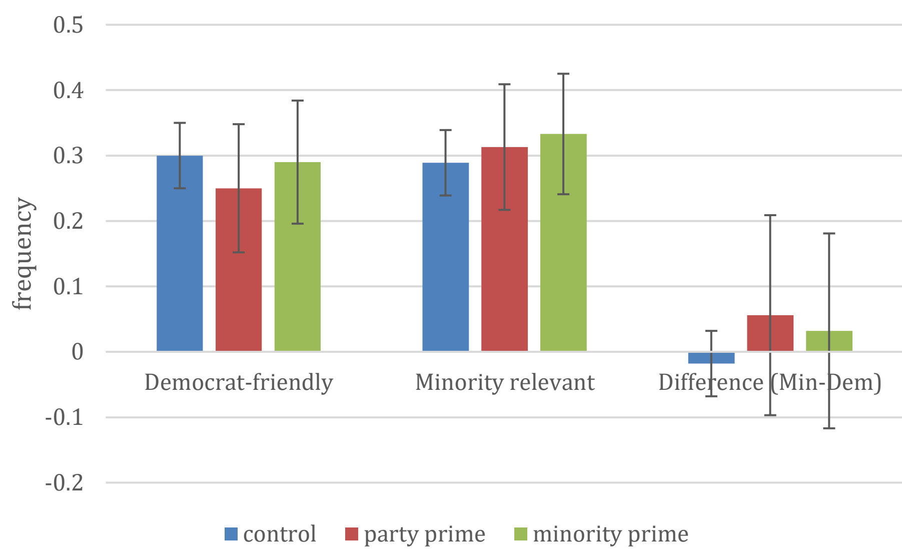

Frequency with which non-white respondents said they were most likely to *read* the story favorable to their party or relevant to their minority group after being primed with their party identity, their minority identity, or with no prime (control). When priming identity, we do not see the same effects on selective exposure as we do on selective curation in Figure 49.

Results: Priming Party and American Identity

In the experiment with white subjects the results are both inconclusive and underpowered (N=103). I report them here nonetheless. As with the results in previous studies, Figure 51 shows the frequency with which partisan respondents said they would choose the story that reflects well on their party or the other party, again in three conditions: when subjects had a partisan identity prime, an American identity prime or no prime. As none of the results near statistically significance we cannot draw any inferences here.

Figure 51. *Effects of Priming Identity on Shifting "Imagined Audience" III*

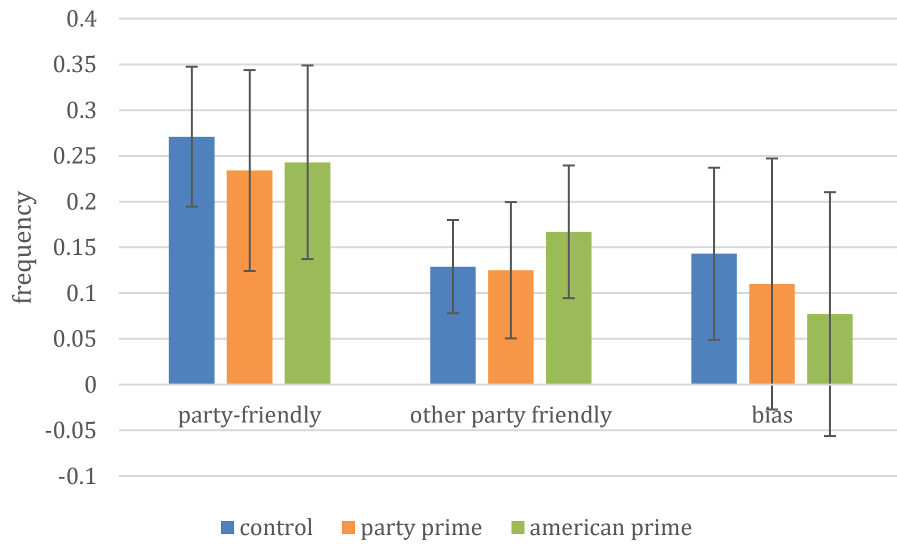

Frequency with which white respondents said they were most likely to share the story favorable to their party ("party-friendly") or the outgroup party ("other party friendly") after being primed with their party identity, American identity, or with no prime (control).

Expectation 3: We are motivated to share stories favorable to our group in order to secure our inclusion in that group

Finally, we arrive at the question of motivation. So far we have seen that social media users, as commonly thought, do tend to share information that would endear them to their political social groups. We also have - tepid - evidence that social media users have those social groups in mind as they select information to share. But we have not yet directly hit on what motivates those users to share stories that glorify their ingroup while damning the outgroup. Here I test the proposition that social media users' curatorial choices are motivated in part by the need to solidify their inclusion in their social groups.

To test that hypothesis, I introduce a manipulation into the experimental design to variably decrease or increase subjects' need for inclusion, observing if they then, indeed, change their willingness to share stories that signal their partisan colors.

For that manipulation I lean again on work in identity and social psychology and use the "self-affirmation" treatment designed by Claude Steele (1988). Self-affirmation theory, as discussed above, posits that humans have a need to perceive themselves as good and competent, and that when they experience a blow to that perception they are motivated to re-establish a sense of self-worth. In Steele's self-affirmation treatment subjects "affirm" their sense of self by contemplating a time when they exhibited qualities that they value; once subjects so affirm themselves, they have less of a need to find other ways to bolster their sense of self. Although first used to support a counter hypothesis to Festinger's cognitive dissonance theory, it has since been used in experiments in other contexts, including selective exposure (e.g. Cohen et al, 2007).

In using a "self-affirmation" treatment, I again make a conceptual leap from the internal self to the external social self. In Steele's conception, humans need to maintain a perception of themselves as good and capable. But if we take the Kurzban view, as discussed earlier, that internal self is merely am image used to project to one's social connections; we need to see ourselves as good and competent in order to project that socially fit image. That internal image also then functions as an internal social meter, tracking one's social worthiness; if our experience of self-worth is high, we will likewise experience that we are in good social stead, but if our self-worth dips it will be a signal that our social standing is at risk. With this Kurzban lens we can then hypothesize that those who have a strongly affirmed self will feel assured of their social image and have less need to signal their group bona fides. Conversely, those who are less affirmed will have a greater need to bolster their inclusion in their social groups, and so will be more apt to want to signal group allegiance.[70]

As with earlier studies, I recruited Mechanical Turk workers who had previously indicated their partisan identity to take a "10 minute survey." The outcome instrument for this experiment was the same as in most of the previous studies; subjects were presented with sets of five recent news stories and asked to place in order which studies they would be most likely to share on social

---

[70] At least one study, that shows that marginal group members are most likely to express derogation of the outgroup, supports the view that expressions of group solidarity come from a need to strengthen inclusion (Noel et al, 1995).

media. In this study, however, subjects were randomly assigned to either receive a self-affirmation treatment, to be in a control group or, in a second version of this study, to be in a “de-affirmation” treatment. In the affirmation treatment, subjects were first asked to place a set of values in the order most important to them and then to write about a time when they felt they exhibited the most important value. In the de-affirmation condition, subjects were instead asked to write about a time they failed to exhibit that value. Subjects in the control condition wrote about a movie they recently saw.

The study was run on two separate occasions, first in January 2017 and then again in June 2017. The first round, in which 349 subjects participated, included the affirmation treatment and control conditions. Results were in the direction expected, yet underpowered. I ran the experiment a second time, introducing the third - de-affirmation - condition, to see if I could strengthen the manipulation. Over the two rounds, 798 subjects completed the survey, 49% Democrat and 51% Republican.

I present the combined results of the two experiments for the control and affirmation conditions in Figure 52 (since both conditions were present in both studies). There we can see weak evidence for the hypothesis; partisans are less apt to select the partisan story if they have been affirmed, yet the difference doesn’t meet standards for statistical significance (p=0.14) in spite of the experiment being sufficiently powered (N=668). (Adding a de-affirmation condition did not, as expected, strengthen the manipulation.)

Figure 52. *Affirming Away a Need for Inclusion*

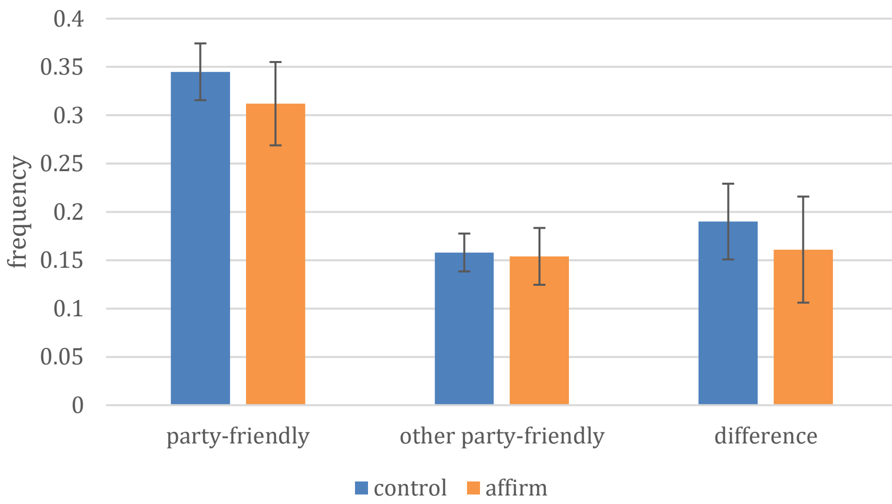

Frequency with which respondents said they were most likely to share the story favorable to their party ("party-friendly") or outgroup party (other party-friendly") after receiving a self-affirmation treatment or with no treatment (control).

What might explain the weak results in this experiment, other than the possibility that the hypothesis is wrong? Again, we cannot attribute the weak results to the experiment being poorly powered; if we do not see an effect with 668 subjects in two conditions there is unlikely an effect worth noticing. Nor is the instrument likely at fault; Steele's self-affirmation treatment has been widely used tested and we have evidence from half a dozen earlier experiments that our outcome variables are strong.

One possible explanation for the weak results are, however, that Steele's intervention does not, indeed, affect one's "need for inclusion"; in other words, the need to perceive oneself as competent and good is not the same as the need for one's group to see one as competent and good, as I posited. If this were the case, then a different instrument that more directly affects the need for inclusion might produce an effect. For example, one could use K.D. Williams social isolation intervention (2007), which induces a feeling of exclusion through an online interaction in which subjects are kept out from a game of catch. There is evidence that experiencing rejection increases identity and connection with one's social groups (Knowles & Gardner, 2008). I avoided using Williams' instrument in part because of the harm (although minimal) it produces and also because it is known to backfire, making some subjects anti-social. It could, however, more directly and successfully induce a higher need for inclusion in subjects. As it stands now, regardless, we do not have strong evidence that it is a need to be included and accepted in a social group that prompts partisans to share stories that favor their party.

Adding Up the Evidence for Group Impression Management Theory

In this chapter, I presented a series of survey studies designed to test Group Impression Management theory, the proposition that social media users share political stories in order to signal their group allegiance and thus secure their inclusion in their political groups. While the results of those experiments are consistent with the theory – and better explain observed behavior than does Selective Exposure theory – we do not find strong evidence to support all elements of the theory. In particular, we have little support for the claim that users are motivated to secure their group inclusion when posting political news on social media.

What the experimental results do demonstrate is that, in an experimental setting, users show a preference for stories that signal group loyalty. Subjects indicated they are more apt to post stories that generally favor their political group, either by confirming the group's beliefs or by extolling the ingroup at the expense of the outgroup. We also have strong evidence that partisans and ideologues prefer to share stories that explicitly tout the ingroup's superiority. To a lesser degree we saw indications that users likewise prefer to share stories that show the outgroup has moral failings and that the ingroup are effective leaders.

The experiments in this chapter also offer evidence – though lukewarm – that users are, indeed, communicating to their political groups when they share political stories. In a group of subjects that hold dual social identities – with their political party and with their racial minority group – we saw that their selection of what story to share shifted when primed with one identity over the other. This shift suggests that users do, indeed, have an "imagined audience" in mind when sharing on social media, an audience that mirrors one of their social groups.

But while users may be speaking to their political groups and they clearly have a preference to share stories that would signal group allegiance, we do not have direct evidence that they share those stories in order to strengthen their inclusion in those groups. The crux of Group Impression Management theory – that users are motivated to secure their acceptance in a group – finds only weak support. There is circumstantial evidence, but no smoking gun. In the conclusion I propose ways forward both for re-testing that conjecture and exploring other motivations for why users share political news.

Conclusion

When I began working on this dissertation in 2014, the original plan was to study how - and how much - social media solidifies our political echo chambers. My initial inspiration for studying political behavior, indeed, had been Cass Sunstein's writings on the dangers of echo chambers online and in civil society. Sunstein's work and my growing concerns about social media and polarization, in turn, led me to work for Eli Pariser as a research assistant on his concept-conceiving book *The Filter Bubble*. By the time I arrived in graduate school a year later it was almost a truism - to me at least - that social media was narrowing our worldview, protecting us from dissonant voices and endangering democratic society. My research, the plan was, would both serve as further confirmation that social media strengthens our echo chambers and hopefully give some insight into how to pry those chambers open.

But as I began to dig into and examine the mechanisms of social media - modeling diffusion patterns, examining Twitter users' retweeting choices, and comparing levels of homophily on and offline - I kept bumping into the same dis-confirming evidence. Social media did not appear to be filtering out challenging information; if anything, it was creating cracks in those chambers, exposing users to information they would less likely come across off social media. When I compared levels of political homophily on social media and in real life I found that, contrary to common thought, our friends and acquaintances on social media reflect our own political colors no more than do our friends offline. Even if we were to surround ourselves with friends who are ideological mirrors of ourselves, those online friends would be apt to send us political information from across the political divide; among the group of Twitter users followed in this work, even the strongest ideologues shared a sizable amount - 20% - of stories from counter-ideological news sources. Finally, modelling diffusion in polarized networks made it apparent that social media sites don't further filter our information for us, as intuition might have it - rather they lead to a distribution of information that is more diverse than the friends we connect with.

Fortunately - from the perspective of any self-interested researcher who abhors a null result - as I saw my original thesis crumble one study at a time, my data and models repeatedly pointed to another - unanticipated - phenomenon. Although social media appeared to expose users to information from the opposing political camp, the stories that percolate to the top tended to be articles from each side's extremes. As seen in the Twitter data I collected as well as data from other scholars, what users reshare is more often than not stories from the DailyKos's and Breitbarts of the online news ecosphere. Balanced, nuanced stories from NPR and the Economist, in contrast, lose out in the battle of re-tweets. Likewise, when it comes to the spread of those "extreme" stories - stories that are dogmatic, tribal and emotionally charged - diffusion is not a potential dampener as it is with echo chambers; rather diffusion amplifies the reach and dominance of extreme news.

My academic journey - from conviction that social media solidifies our echo chambers to the realization that it is more likely a breeding ground for extreme information - reflects in large degree public perception of social media's political role over the same time period. In 2014, when my research began, "echo chambers" and "filter bubbles" were known and well documented phenomenon. Eli Pariser's 2011 *The Filter Bubble* had invented a term that had already entered common parlance in discussions about online media and society.

Figure 53. *Filter Bubbles in the Public's Mind*

Relative frequency of Google searches for the term "filter bubble" 2014 until today. (From Google Trends)

Academics likewise produced a steady stream of studies that analyzed and visualized our online bubbles, each with accompanying visuals of the red and blue "hair balls" of our online networks.

Figure 54. *A Popular Visualization of Online Echo Chambers*

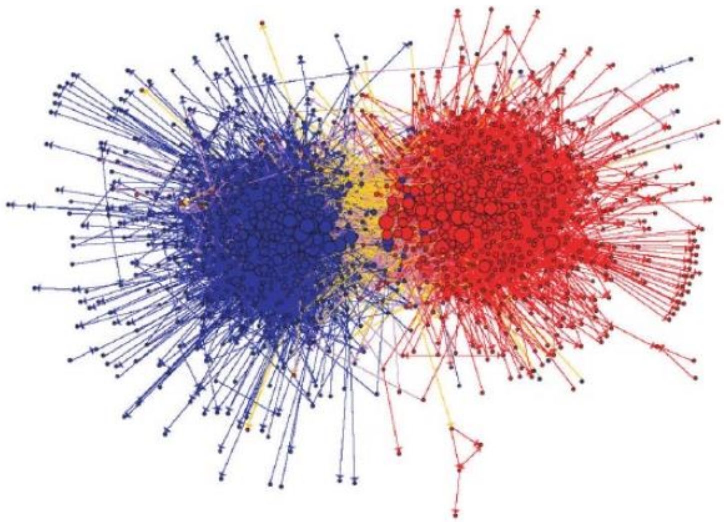

Image of blogosphere network from Adamic & Glance (2005).

Echo chambers, true, were not the only social ills associated with our online lives in 2014. Trolling, flame wars, bullying - and other versions of uncivil or harassing behavior - were also well recognized scourges. But while platforms and academics worked to figure out how to minimize online harassment, by and large these antisocial behaviors weren't seen as a threat to political society.

Those threats did materialize in 2016. A new set of trolls - this time savvy at manipulation and often bent on disrupting democracy - emerged in the presidential election that year. "Fake news" sites proliferated, taking advantage of social media to push news stories that fed into the fears and fantasies of left and right ideologues. Foreign actors set up bots and fake Facebook groups to further promote and disseminate the spread of stories designed to stoke inter-group animosity. With less of an intention to harm the democratic process, but perhaps with equally harmful consequences, hyper-partisan news creators and alt-right groups likewise capitalized on social media's capacity to drive traffic to sensationalist stories and extreme views. All those trends were part of a larger phenomenon of what I term the rise of "extreme news" on social media. Though it comes in many forms, is sometimes fake but more often well-spun reality, extreme stories share three traits; they dogmatically uphold partisan beliefs, use emotion rather than reason to engage users and stoke tribal disagreement and conflict.

Figure 55. *Fake News Enters the Discussion*

Relative frequency of Google searches for the term “fake news” 2014 until today. (From Google Trends)

Today, social media is still seen as a creator of cozy echo chambers that coddle Americans into thinking everyone shares their views, but also as a conduit and amplifier of conspiracy theories, outrage politics and ideological rhetoric wars. We are left with twin demons of online polarization - protected in the belief that we are right in our political views, on the one hand, while also having our fears stoked by extreme news that portrays the other side is corrupt and dangerous.

Except, as this dissertation helps demonstrate, that portrait of social media is not completely accurate. Social media users do tend to surround themselves with like-minded friends on social media and see more information that reflects their political beliefs, but our social media echo chambers are also quite porous. That view of social media is not just supported by this dissertation. Since 2014, other scholars have likewise poked holes in the information bubble narrative. For the most part online news consumers are exposed to a healthy diet of centrist, mainstream news with considerable overlap between right and left readers (Guess, 2015). Social media does lead users to more ideologically congruent news sources than do search engines, but also to a more diverse array than if those users were to go directly to their preferred news sites (Flaxman et al, 2016). Looking at data from Facebook and Twitter, it becomes apparent why that is the case; social media users connect to a fair number of counter-ideologues (Bakshy et al, 2015, supplemental; Faris et al, 2017).

But even though claims that social media reinforce echo chambers are misinformed, that does not mean platforms and scholars should be complacent and not put effort into breaking down the walls of our echo chambers online even further.

Indeed, social media offers an opportunity to bridge left and right that does not exist offline. In the real world, connecting liberals and conservatives faces geographical barriers. Though stark maps of "blue vs. red states" often exaggerate the physical distance between left and right, Americans who live in pockets of liberalism (Hello Berkeley) and conservatism (Hi Colorado Springs) would find it difficult to find a counter-ideologue to have coffee with even if they were inspired to reach across the ideological divide. (When this liberal tried to find conservatives in Manhattan once upon a time, she had to go to a Young Republicans meetup, and even there only found economic conservatives). On social media, however, there are no geographical barriers. An Oakland liberal can converse easily with an Oklahoma City conservative.

Social media likewise has leverage in altering our information sharing behavior that is limited offline. Changing human behavior is no mean task. It is possible, however, to "nudge" behavior in one direction or another by altering the environment a person is in. Altering Americans' real world environments would, of course, be a monumental - and foolish - enterprise. On social media, however, modifying a user's environment is a matter of changing some code. The setting of a social media platform is completely under the control of its designers; it opens the possibility of shifting our behavior in ways that not only appeal to social designers but possibly also to users on that platform.[71]

---

[71] Here I am adopting the Thaler & Sunstein (2009) view that humans by and large want to be their “better”, “forward thinking” selves and so nudging them from being their “immediate gratification” selves is not inherently manipulative.

How could social media platforms help make our echo chambers even more porous? This dissertation does not offer any silver bullets, but it does provide some clues. For one, the models in Part I can help direct platforms where to focus their design energies. When looking at how to diversify users' information environments social media presents two levers: it could help create more connections between left and right users or it could nudge those users to share more diverse information. (It could also, of course, more directly tweak its algorithm to promote more diverse stories or simply place diverse news in ads or on users' feeds.) In choosing which lever to capitalize on, platforms may want to consider which will get more miles for their money. In Chapter 2, we saw that all nudges do not produce the same size effects. As information diffuses through polarized networks the level of information homogeneity ("information sorting") moves towards an equilibrium. That equilibrium, as shown in Figure 56, depends on how homophilous a network is and how biased users are in what they share.

Figure 56*. Information Sorting Equilibria: Reprise*

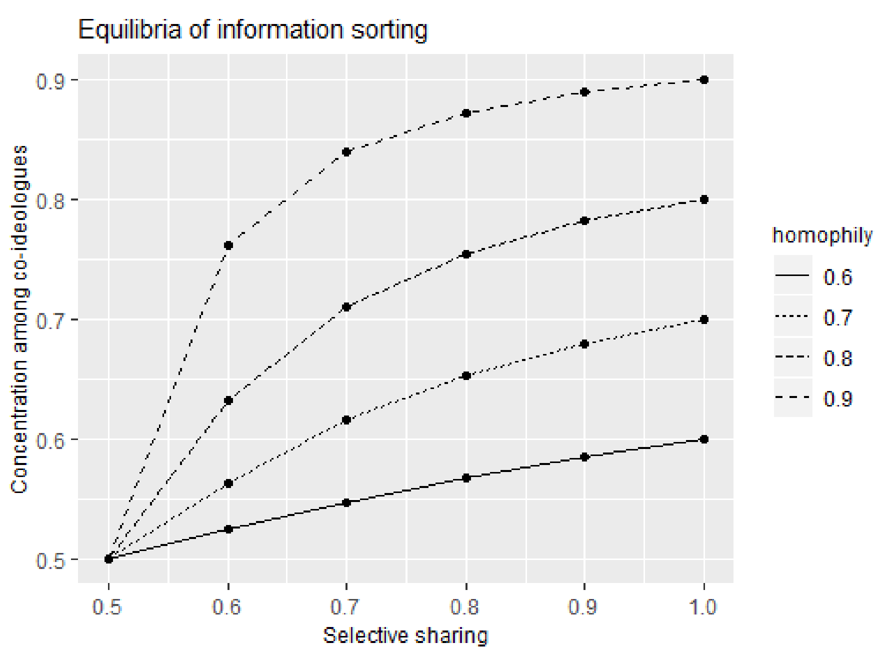

Information sorting equilibria ("concentration among co-ideologues") reached in networks of varying levels of homophily and selective curation ("selective sharing"). Based on simulations in an infinite network.

The graph illustrates that the relationship between information sorting, homophily and selective curation is not a linear pay-off. Depending on its starting point, a platform will get different gains whether choosing to nudge its users toward greater homophily or less biased sharing. If users are highly ideological in choosing information to share (with a curation bias greater than 80%), pushing them to be less biased will have a small effect on diversifying users' information environments. At high levels of biased curation, then, more will be gained by nudging users to have less homophilous friend circles. On the other end, if users are already fairly unbiased in what they forward to their friends (with a curation bias below 80%), any move toward less bias will have a relatively large impact on the network's information diversity.

What we saw in examining the sharing behavior of social media users on Twitter suggests, however, that users are already fairly unbiased in what they choose to re-share. Indeed, this work's sample of Twitter users were on average completely unbiased; users tended to share the same proportion of liberal and conservative news that they were exposed to themselves. If those Tweeters are anywhere near representative of the average social media user, there may not be much room to encourage social media users to be less ideologically biased in what they reshare; they are already equal opportunity sharers. The only way for platforms to break down echo chambers any more would be to nudge users to build more connections to counter-ideologues.

This dissertation doesn't propose ways to build those bridges, but it does ask: is that something we'd really want to do? Bubble bursting advocates like Cass Sunstein argue ideologues should be exposed to views from their opponents with the hope that such exposure will lead to, if not consensus, then at least a level of mutual understanding and tolerance that sets the foundation for politically necessary compromise. Such an outcome, however, presupposes that the information partisans are exposed to will inspire understanding and tolerance. We have good reason, however, to suspect that is not the case on social media. Cross-ideological exposure *can* lead to greater understanding and tolerance, but as practitioners and theorists of deliberative democracy know, if interaction between contentious political groups isn't "appropriately empathic, egalitarian, open-minded, and reason-centered," the benefits of cross-ideological engagement may go missing (Mendelberg, 2002). If a baseline of facts is missing or the rules for discussion aren't clear, any bridge that is created will end up in flames. Flame wars are, indeed, what too often materialize when liberals and conservatives meet online.

Similarly we should expect that a news article posted by a counter-ideological friend on social media would do little to inspire political comity. Sampling recent stories from one of my conservative friends on social media, as illustration, I see the following posts: a news report claiming that climate scientists fake data, an op-ed claiming that the most efficient states are Republican led, and a post touting the accomplishments of Trump. As a liberal reading these posts, I confess, I do not feel I have a greater understanding of the world nor am I reassured that conservatives have well-founded reasons for their views. I might even be tempted to think this friend - and other conservatives - are biased partisans with conspiratorial tendencies and not very smart. (When I speak with my friend, in contrast, I know this not to be true; he is actually intelligent and thoughtful.)

I am not alone in having this reaction to seeing conservative posts. Jaime Settle (2019) examines how users form political attitudes on Facebook and argues that the Facebook feed is "uniquely suited to facilitate processes of polarization." Users make inferences about counter-partisans when they see their political posts online. More often than not those inferences serve to increase negative stereotypes about those who do not share their political views. A similar process may explain why, when researchers induced Twitter users to follow the account of a counter-partisan for a month, those users did not increase their appreciation of their political opponents but rather became further entrenched in their partisan views (Bail et al, 2018).

The work in this dissertation gives insight into why exposure to counter-partisans may lead to polarization. The inferences we make about others depends in part on what they are sharing. As we saw among the Twitter users we followed - and as seen by other researchers - the problem is not that users are posting news with a left or right lean, but that they are sharing stories with *extreme* left or right leans. Diffusion models, moreover, show that those individual biases will amplify the reach of extreme tweets - relative to more nuanced and balanced news stories - as tweets move through a homophilous network.

In sum, breaking through our echo chambers on social media may do little to reduce polarization. Instead, building online bridges may have the unintended - and harmful - effect of pulling left and right farther apart. If we want to mitigate social media's capacity to fuel polarization, then, bursting bubbles is not enough - it will be crucial to stem the spread of extreme news.

Social media platforms are already working to rein in the worst offenders of extreme media - fake news and hate speech. They do so with brute force moderation, banning users that propagate hate speech or are known purveyors of fake news (Brownlee, 2019; Constine, 2018). But much - or most - extreme news stories are neither fake nor explicitly hate-directed. They are, however, as capable of inspiring outrage, animosity and self-righteousness as a fabricated conspiracy theory or blatant bigotry.

If platforms were to focus on mitigating the spread of extreme news more broadly, how might they do so? For starters, they may not want to use blanket censorship. Culling extreme stories comes with at least two obstacles. First a platform would need to create a reliable algorithm for identifying an extreme story; even with the strides being made in AI as of the writing of this dissertation, such an algorithm (machine learned or otherwise) seems far off. More dauntingly, deploying such censorship would likely face backlash among users who will not be pleased when they see their OccupyDemocrat and Fox stories disappear.

Algorithms could be deployed to, instead of outright censor, simply downgrade extreme stories, giving them less prominence in friends' feeds. Google search uses such a tactic, rewarding news sources it deems high in expertise and trustworthiness (The Economist, 2019); if social media sites were not already using such weighted measures it would surely be within their reach to do so. If users are less likely to see extreme news stories they will necessarily be less likely to reshare them. But even such weighting risks cries of censorship.

A different approach is to steer users away from choosing to share incendiary news. As discussed above, social media platforms have full control over the environment in which their users interact; design features can be tweaked to likewise tweak users' behavior. To know what design features might nudge users away from sharing extreme news, it is necessary to know what motivates users to share political stories in the first place. That, of course, was the object of Part

II of this dissertation. I proposed that social media users are driven to project an image of themselves that would make their social groups accept them as a member. If ideologues and partisans are concerned with signaling that they are a member worth having - i.e. one that is true and loyal - we would expect them to share news stories that reaffirm the group's beliefs and generally show the ingroup in a positive light relative to the outgroup. We'd also predict that stories that have the hallmarks of extreme news - i.e. are dogmatic in their confirmation of group beliefs, emotionally laden and tribal (extolling the ingroup while denigrating the outgroup) - do a better job of signaling group allegiance and so would be more shareable yet.

Experiments in Part II confirmed many of those expectations. We did indeed see that when social media users have a choice, they are more apt to say they'd share stories that make their group look good, in particular stories that explicitly proclaim the superiority of their group members. We also saw some evidence - although weak - that users are speaking to their politically relevant groups when they share political stories. What the experiments failed to shed light on, however, was the ultimate motivation for why users share those stories. Experiments intended to find evidence that users are motivated to secure their inclusion in a group were inconclusive.

The studies in Part II, nonetheless, suggest a couple of interventions that might diminish users' predilection to spread extreme news. Far more so, those studies point to how much more work needs to be done to understand what does, indeed, drive our sharing behavior. Below I review elements of the theory I proposed, for each suggesting possible relevant interventions and offering a roadmap for further research.

The Imagined Audience

What We Know and Need to Learn:

In the experiments in Chapter 7, we observed some evidence that users are communicating with their social groups when they post news stories on social media. Democratic Black and Latino exhibited different sharing preferences when cued to think of their minority group vs their partisan group, a finding which suggests the stories users select to share are determined in part by the audience users have in mind when making that choice.

Those effects were not, however, seen with White subjects when primed either with their partisan or American identity. But a lack of an effect in those experiments could be due to their low statistical power. To confirm that a user's imagined audience can be adjusted, more - well powered - studies would need to be conducted across other populations with other dual identities,

Possible Interventions

If users do, indeed, have their partisan and ideological groups in mind when they share political news - and their object is to signal shared beliefs and group loyalty - then shifting users' "imagined audience" should also shift the impression they want to make. Any move away from thinking of their political groups should result in sharing less extreme news. If, for example, a Democrat has only fellow Democrats in mind when sharing news about abortion policy they may be more apt to share a story that strictly upholds a pro-choice position while disparaging those who support limits to abortion. Such an article, while buying cred with their Democratic friends will serve to polarize friends who are not lock step progressives. If, however, they are cued to think of another social group - perhaps their work friends, church community or all Americans - they might be less willing to send such a pure signal of Democratic solidarity and instead choose to share a story that, while still supportive of a pro-choice position, is less strident and dismissive of those who disagree.

There are number of ways a platform might cue a user's non-partisan identities. On platforms in which users subscribe to "groups" a platform could choose to make those other groups more visually salient, either in general or when users click on the "share" button. Similarly when a user clicks to "share," a platform could pull up the profile thumbnail of a variety of the user's friends. Such interventions might serve to remind users that their post will not only be seen by fellow partisans and ideologues, but by a wide array of connections - some of whom might not applaud the user for their extreme story.

Securing Inclusion

What We Know and Need to Learn

This dissertation posited that, in sharing political news, social media users were primarily motivated to secure inclusion in their groups. Experiments, however, failed to find evidence that is the case. I used a self-affirmation manipulation to reduce users' need to seek approval and acceptance from their social groups, yet such a manipulation had little effect (and no statistically significant effect) in deterring users from sharing group-affirming stories. To see if users are, indeed, motivated to signal group solidarity further experiments are necessary. As discussed in Chapter 7, it may be that the manipulation used did not, indeed, affect a user's need for group inclusion; an alternative manipulation that could be tried is William's (2007) social isolation intervention which more directly makes subjects feel socially rejected and thus, possibly more motivated to ingratiate themselves with their social groups.

Possible interventions

If further studies show that users are motivated to secure group inclusion, platforms might nudge users away from sharing extreme stories by re-affirming their social standing. One way to do so would be to remind users of their social acceptance, for example showing how many friends have liked their posts. A risk here, however, is that showing a tally of likes could gamify their behavior, incentivizing users to push that tally up. If they recall that their more extreme posts get the most responses (as is likely) this could spur them to post more incendiary stories.

Increasing Status

What We Know and Need to Learn

The theory of motivation I present posits that users are not only driven to secure inclusion in their groups but also to raise their status within those groups. I did not test this proposition directly, but some exploratory work I did on traits that predict ideological and partisan sharing indicates that status may be a driver. In that exploratory work, described in Appendix III, I asked subjects to say how much they agreed with a series of statements including "I tend to take charge." While I expected this statement to act as a proxy for a need for status and so to increase partisan sharing, I found the opposite; those who rated themselves high in leadership were actually less biased in the news they shared. In contrast, another group of subjects - who tended to agree with the statement "I value my intelligence as an important part of me" - were the most partisan sharers.

Those twin results - one expected and one unexpected - suggest that users may post stories to impress others - yet not only in the way I assumed. It is not news that partisans who are most politically savvy are also those who hold their beliefs most firmly (Zaller, 1992). Those more certain of their outlook are also more likely to be politically engaged (Mutz, 2002) - and thus the most prolific political news curators on social media. Sharing political stories is an opportunity to demonstrate their intelligence and so raise their status. (Or, at least, they imagine it will raise their status; how many of their friends admire them or simply roll their eyes is up for debate.) But, whereas I posited that users are motivated to raise their status within their social groups, this desire to look smart quite possibly transcends group allegiance. It could be, rather, that these users are both motivated to share political stories and happen to have strong partisan and ideological beliefs; it follows they will be likely to share extreme stories in general, not just to gain status with a particular group.

The other finding regarding subjects who see themselves as leaders could also indicate that status-seeking is a motivator for sharing extreme news. Whereas I anticipated that those who see themselves as leaders are status seekers, however, the opposite may be the case; leaders might be those who see themselves already as having high status. If that is the case they don't need to jockey for position by signaling they are even more devout believers; they may even be in a position to challenge their group on their beliefs - which is what these subjects seem to do.

Possible Interventions

If an insecurity about one's status does drive partisans to share more firmly group-affirming - and so extreme - stories, then one tactic to mitigate that drive would be, as above, reassuring users that their status is already high. Yet, as with the caution in mitigating the drive to be included, such a signal risks incentivizing status-seeking behavior. The desire to impress others may be so strong, moreover, that it is foolish to imagine that a design feature could make users more secure in their status.

A better approach might be to alter users' perceptions of what kinds of stories are more likely to raise their status. Users who want to signal their smarts, in particular, would be sensitive to being caught out sharing an article that comes from a source known to be careless with its facts and analyses. Facebook does this to an extent already, linking news stories to Wikipedia pages with information about their source. Yet, although users could read the full Wikipedia entry to figure out if the source has a strong reputation, that demands work. Platforms could instead link to independent ratings of that source's reputation for reliability and sound analysis. If such information is made salient, users who value their own reputation for intelligence might be more careful not to share stories from extreme news sources.

Threat and Fear

What We Know and Need to Learn

Another element of the theory of motivation that I did not test is that users are driven to share extreme stories when they feel they - or their ingroup - is threatened. Sharing an extreme story - say about a corrupt leader, a policy that threatens the group's values, or an event that is deemed an existential threat - is a way to alarm the group and galvanize them to act in unison to combat the threat. There are a number of ways to test this proposition - by using manipulations that either mitigate or induce a subject’s baseline feeling of threat. Subjects who are made to feel more generally safe and secure should be less likely to share alarmist - extreme - stories.

Possible Interventions

Any intervention that decreases a user's general experience of threat might likewise diminish their motivation to share alarmist stories. One approach would be to mitigate the experience of threat from what most immediately precipitates it - that is, the news story the user sees on their social media feed. When reported on by an extreme source a threatening news event may indeed seem more alarming yet. A less sensational source is more likely to put an event in context and give a more balanced, nuanced picture of its relevance. An intervention then might be to just offer "other sources on this story" when a user sees an extreme news article. Whereas the Breitbart or DailyKos story might stir fear, a New York Times or even Fox News story may offer a take that is less alarming.

Far More Work to be Done

The suggested interventions and roadmap for further research above are by no means exhaustive. The theoretical framework for understanding sharing behavior I offer in this dissertation, likewise, is merely offered as a starting point to think deeply and broadly about what drives our choices to share political information online social media.

Hopefully I have convinced a reader or two that this is an important question to ask. It is impossible to predict what our media environment will look like fifty - or even ten - years, but in the immediate future we can expect that social media users - that is, all of us - will in large part determine the information diets of our fellow citizens. What makes up those diets will influence how well civil society can work together to address national and global problems. As I have argued in this dissertation, we should be concerned about the amount of extreme information that funnels through our news feeds; by heightening tribal divisions between left and right extreme news inhibits our ability to see common ground and, more ominously, raises animosity between political tribes. What work other scholars do to understand what drives our propensity to share extreme news - and thus be able to discover ways to mitigate that drive - will help diffuse inter-party animus and, I hope, help strengthen the fabric of civic society.

Appendices

Appendix I: Agent-Based Simulations Across Homophily Levels

In Chapter 2, I presented the results of diffusion simulations using agent-based models on graphs with 0.8 homophily. Here I present those results across graphs with levels of homophily from 0.5 to 0.95.

Figure 57. *Diffusion in Agent-Based Models: Information Sorting at the Network Level*

A. Binary Random

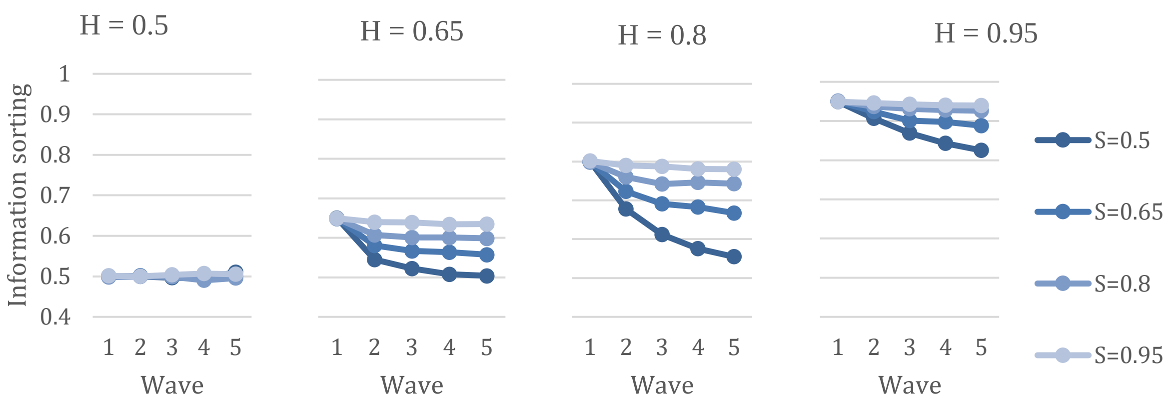

B. Binary Small World

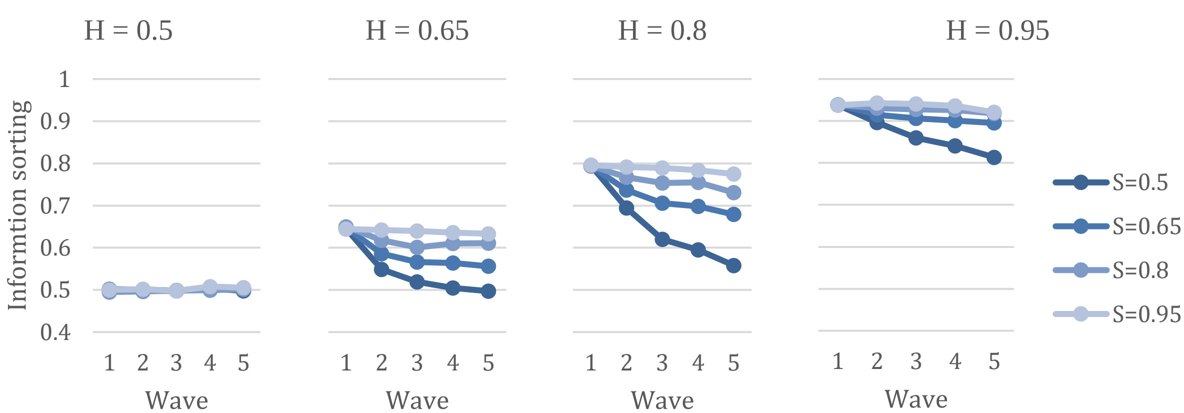

C. Random Continuous

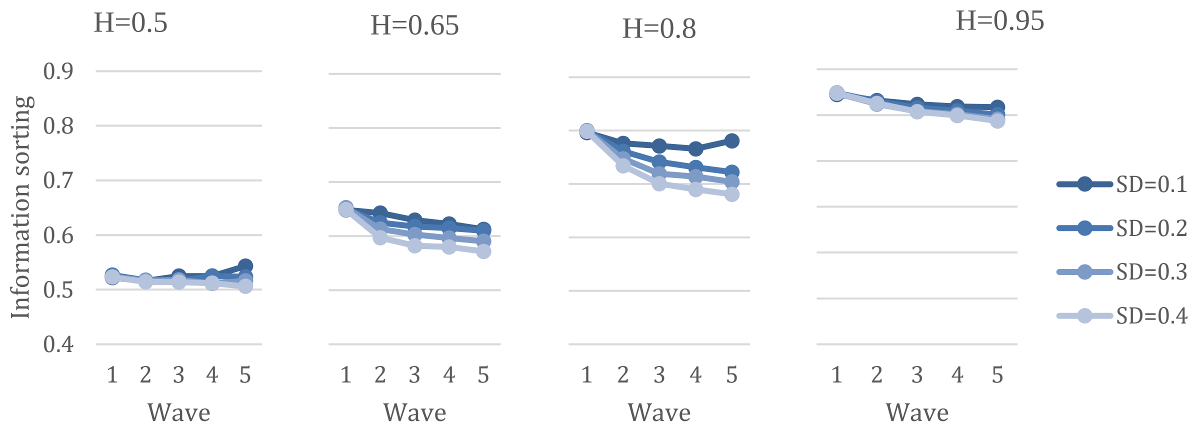

Proportion of nodes exposed to congruent messages ("information sorting") over waves of diffusion for different levels of selective curation (S for binary networks and SD for continuous networks), in simulated diffusions on a) binary random networks, b) binary small world networks and c) continuous random networks. Initial messages share ideology of seeded node. Looking across networks with homophily (H) levels of 0.5, 0.65, 0.8, and 0.95. All networks have average degree of 8.

Figure 58. *Diffusion in Agent-Based Models: Information Sorting at the Node Level*

A. Binary Random

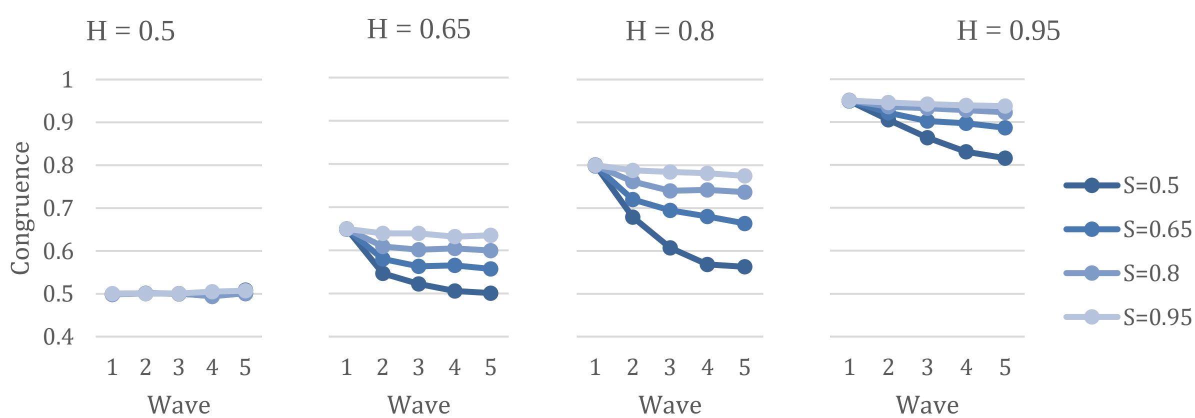

B. Binary Small World

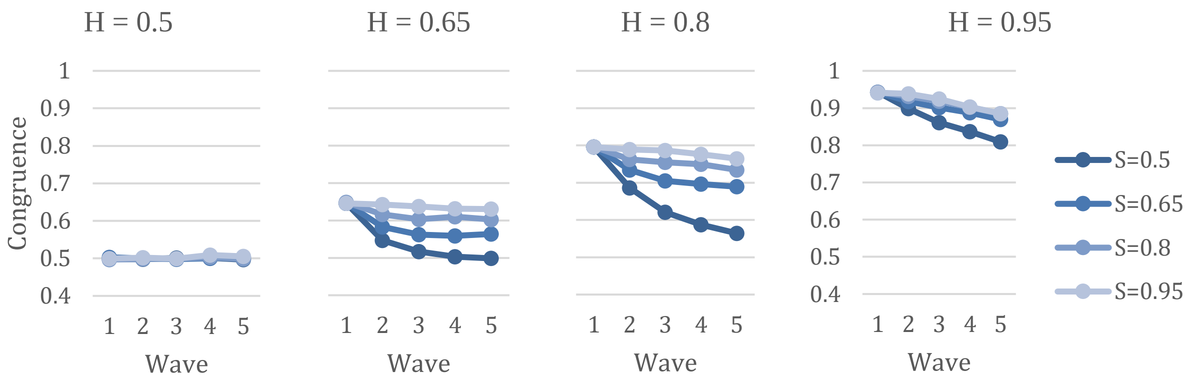

C. Random Continuous

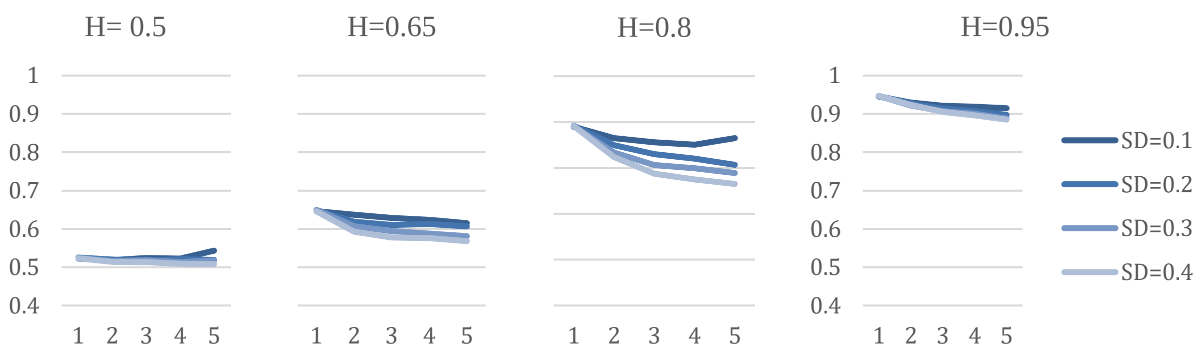

Proportion of messages nodes are exposed to that are congruent ("congruence") over waves of diffusion for different levels of selective curation (S for binary networks and SD for continuous networks), in simulated diffusions on a) binary random networks, b) binary small world networks and c) continuous random networks. Initial messages share ideology of seeded node. Looking across networks with homophily (H) levels of 0.5, 0.65, 0.8, and 0.95. All networks have average degree of 8.

Appendix II: Self-report and Naïve Theories of Selective Curation

An approach scholars often take to understand social media users' motivations is to simply ask them, using self-report surveys. In this dissertation I avoid using surveys for the reason that humans are by and large (if not entirely) unaware of their motivations (Bargh & Chartrand, 1999). The goals we have for sharing information online are, just like any other goals, most likely unknown to us. What explanations we *do* offer for our behavior are usually confabulations or rationalizations we create in the moment to explain our behavior in a way that is both plausible and that creates a positive image of ourselves (Nisbett & Wilson, 1977). I suspect this to be even more true in the realm of political behavior, in which individuals may be particularly prone to rationalize their beliefs and behaviors and be desirous to appear as good citizens.

As a demonstration of why we may not want to trust self-report to motivations, I ran a survey asking a number of social media users, once again recruited from Amazon's Mechanical Turk, why they post political information. I also, in the same survey, asked a subsection of respondents why they think *others* post political information. I presented 201 subjects, a subset of whom said they at least occasionally posted political information on social media (N=89), with ten possible reasons for sharing political information[72]. Among the ten reasons, four were attributed to self-serving or negative motivations, four presented more altruistic or positive motivations and two were relatively neutral.

[72] I had distilled those ten reasons from open-ended responses from fifty other Mturkers whom I had asked "If you had to say, why do you think [people/you] post information on social network sites about politics?"

Table 3*. Reasons for Sharing Political Information*

Self-serving or negative motivations	To validate their beliefs.
	To persuade their friends to believe something that is in their self-interest (that is, the self-interest of the poster, not their friends').
	To get attention.
	To appear intelligent.
Neutral motivations	To find others who think the way they do.
	To show their friends they share their friends' beliefs and values.
Altruistic / positive motivations	To inform their friends about important issues.
	To find out their friends' views on an issue.
	To keep their friends up to date on topics they know their friends care about.
	To engage in discussions where everyone learns from different perspectives.

Ten possible motivations for sharing political information on social media.

Among the subjects who said they sometimes post political stories, I asked a random subset (N=44) “What are the most likely reasons *you* post information about politics or social issues?” (emphasis added here) instructing them to place the ten reasons in order from most to least likely. To the other 157 subjects (including the other social media users who do post political information themselves) I asked “What are the most likely reasons *people* post information about politics or social issues?” (again, emphasis added).

Figure 58 shows the average placement subjects gave for the ten reasons from most likely (1^{st} place) to least likely (10^{th}). When asked about other people, there was not strong consensus (average placements range from 4^{th} to 7^{th} place), yet overall subjects tended to say that users have self-serving motivations in posting about politics; the top two reasons they gave were that users wanted to "validate their beliefs" or "persuade their friends" to believe something in their (the poster's) self-interest. Least likely, according to subjects, is that posters want to keep their friends up to date on issues their friends care about or that they want to learn about their friends' views.

When we ask subjects about their own motivations for posting about political issues, in contrast, there is both more consistency in how they assess their motivations and consensus that their motivations are selfless. Almost all agree (37 out of 44) that one of their top motivations is to "inform their friends about important issues." And few (6 out of 44) are willing to say that one their top motivations is to "get attention." (Rather 24 out of the 44 place "get attention" in 10th place.)

Figure 59. *Reasons for Posting Political Information, Ordered from Top to Bottom*

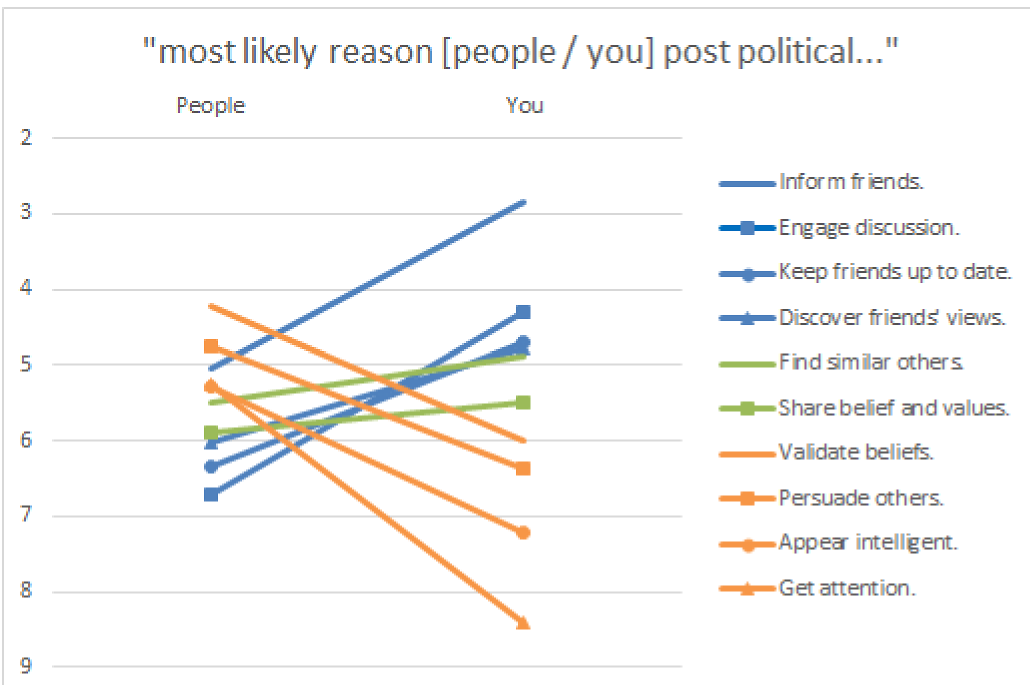

Average order of reasons subjects say they ("You") or others ("People") are most likely (1) to least likely (10) to post information about political or social issues.

The disparity between the motivations subjects attribute to others' behavior versus their own behavior will not be surprising to survey methodologists or social psychologists. As discussed in this work, survey participants are well known to give “socially desirable” answers on questionnaires; having self-serving motivations is frowned upon so subjects will be loath to admit such selfish behavior to a survey giver, even a remote one online. But subjects may not just be trying to fool a researcher; they are likely also fooling themselves. Harkening back to Steele (1988), we know that humans need to think of themselves as good and competent individuals. Since we do not have access to see our true motivations, we will confabulate explanations that put us in a positive light. When a survey participant says she shares information “to inform her friends about important issues” and that she is not motivated to “get attention,” she actually believes her own spin (Kurzban & Aktipis, 2007; Millham & Kellogg, 1980). When it comes to the behavior of others, however, she can be more objective and thus able to think like an intuitive psychologist, correctly divining the motivations of others.[73] Another possibility is that she attributes selfish motivations to others because it makes herself feel virtuous by comparison. Regardless of our objective ability to guess the motivations of others, when it comes to our own self-assessment, it is most likely that a need to perceive ourselves as decent humans is behind our explanations for our motivations.

The reader might propose an alternative explanation; perhaps it is the case that people who post political information *do* have more insight into their motivations and are able to accurately identify those motivations. The disparity between their self-assessment and the assessment of others may be explained by the fact that those who don’t post political information cannot draw on their own experience and self-reflect, so instead must guess about the motivations of others. We can see if that is the case by restricting our results to only the subjects who say they post about politics on social media; in theory they should be able to draw upon insights into their own motivations to inform how others might likewise behave. If we restrict ourselves to that sample, however, we see similar results as above:

[73] Although our assessment of others based on their behavior is far from perfect (Ross, 1977).

Figure 60*. Reasons for Posting Political Information, Ordered from Top to Bottom, Among those who Post Political Information Themselves*

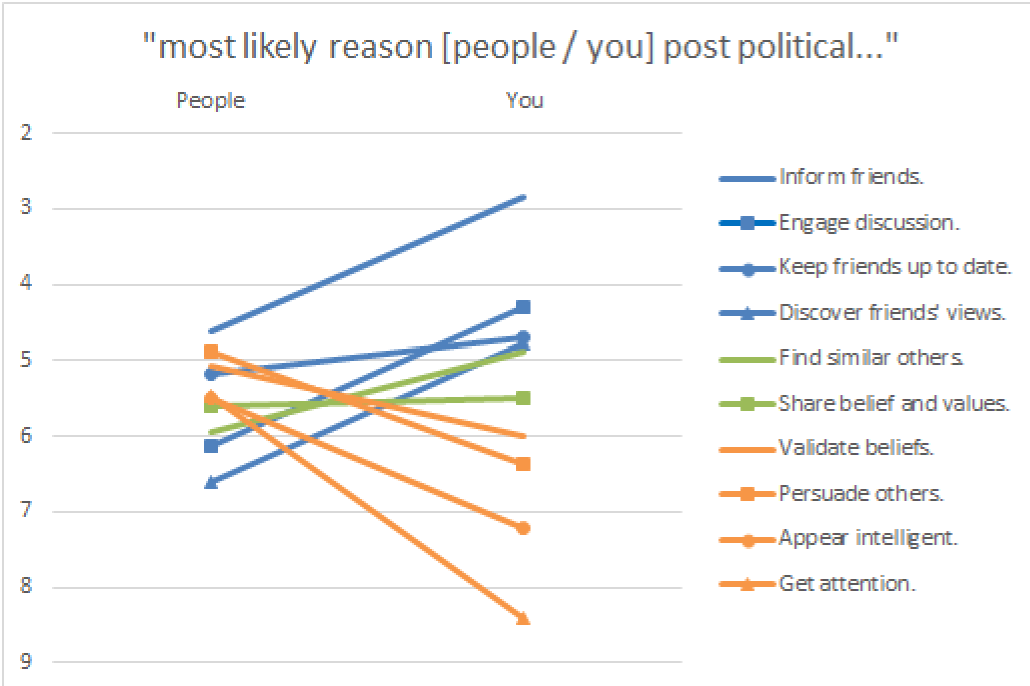

Average order of reasons subjects *who post political information* say they ("You") or others ("People") are most likely (1) to least likely (10) to post information about political or social issues.

Subjects who post political information are somewhat more generous in guessing the motivations of others (they now agree that the most likely motivation is to inform their friends), but they still by and large are more apt to attribute self-serving motivations to others than to themselves. In this case it is harder to make a case for why subjects would see their own sharing in such a positive light while seeing others' sharing so negatively - other than to say they have a self-serving bias. We do not know, however, if that self-serving bias pumps up the estimation of their own motivations or merely depreciates the motivations of others.

The key point is "we don't know." Not knowing how individuals' biases distort their perception of their own motivations, we cannot trust their self-reports. Instead of asking subjects why they do what they do, we are better off using the tools of social psychology and social science; developing theory and testing those theories in experimental settings.

Appendix III: Traits that Predict Selective Curation

Another way to explore what drives a behavior is to see if individuals with certain traits are more or less prone to that behavior. In the case of ideological selective sharing, if we imagine – as this dissertation proposes – that users share partisan news in order to secure inclusion in their political groups, then we would expect users who generally have a greater need to be included to be more careful to signal their political alignment than other users.

In an exploratory experiment, I tested this prediction along with several others about traits that might lead to greater levels of ideological curation. I look both at traits derived from the theoretical framework presented in this dissertation as well as traits in line with the naïve theories proposed by the subjects surveyed in Appendix II. Those traits and their association with ideological curation are presented below.

Need for Inclusion:

There is no existing measure for an individual's need for inclusion; as proxies I instead used two measures – Need for Belonging and Need to Conform – that, similar to a need for inclusion, home in on an individual's connection to social groups and need to be part of those groups.

Need for Belonging, in Leary's conception, "goes beyond a mere desire to affiliate or socialize to a desire to be accepted, form relationships, and belong to social groups" (Leary, 2013). On a 5-point scale I asked subjects how strongly they agreed or disagreed with the two statements "I have a strong need to belong" (Nichols & Webster, 2013) and "If other people don't seem to accept me, I don't let it bother me" (reverse coded) (Leary, 2013).

For *Need to Conform,* I used two items from a composite index for Conformity / Dependence / Need for approval: "I need the approval of others" and "I don't care what people think of me" (reverse coded).

Need for Affiliation. As an alternative theory to Need for Inclusion, I modified an index from Hill (1987), taking items that tap into two dimensions of affiliation; emotional ("If I feel, unhappy I usually try to be around others to make me feel better") and positive stimulation ("Listening to others and relating to them is one of the most satisfying things to do").

Need for Status.

To test the proposition that those who seek status in a group will also be careful to signal group allegiance, I looked for – and failed to find – a measure for Need for Status. Coming up short I instead used a plausible proxy; using items from a composite measure for Leadership: "I tend to take charge" and "I usually wait for others to take the lead" (reverse coded).

Naïve Theories.

Need for Validation. Among subjects I surveyed (discussed in Appendix II), the most popular answer to the question "Why do others post political information?" was "to have their beliefs validated." There being no measure of "need for validation" I instead used low self-esteem as a substitute, reverse coding Hill's one item self-esteem item "I have high self-esteem" (1987).

Need for Attention. Subjects likewise suspect others post political news to get attention. To measure need for attention I use one item from Hill (1987), "I like to be around people who are impressed with who I am," and reverse code "I feel uncomfortable being the center of attention."

Altruism. A common reason proposed for why individuals share political information with their friends is that they are doing so altruistically. To see if altruism does lead to more or less curation I use two items from a composite altruism scale: "I love to help others" and "I take little time to help others" (reverse coded).

Curious / Need to learn. Others claim they share political information in order spark discussion and learn from their friends. If this is the case we might expect to see those with a stronger desire to learn to have different sharing patters. Again, I take items from a composite index for curiosity / need to learn: "I like to know how things work" and "I am thrilled when I learn something new."

Need for intelligence. Finally to test the naïve theory that users post political information to look smart, I invent a mini-index for "Need for intelligence," with the two items: "I value my intelligence as an important part of me" and "I think it's more important to be good than smart" (reverse coded).

Results

Figure 4 charts what degree – if any – the above traits predict partisan curation (in the initial experiment presented in Chapter 7). The items I used for Need for Inclusion (Need belong and Need conform) have no association with partisan curation. The item I used as a proxy for Need for status (Leadership) has a negative association with partisan curation. The only item that successfully predicts partisan curation – by a wide margin – is Need for Intelligence.

Figure 61. *Association between User Traits and Ideological Curation*

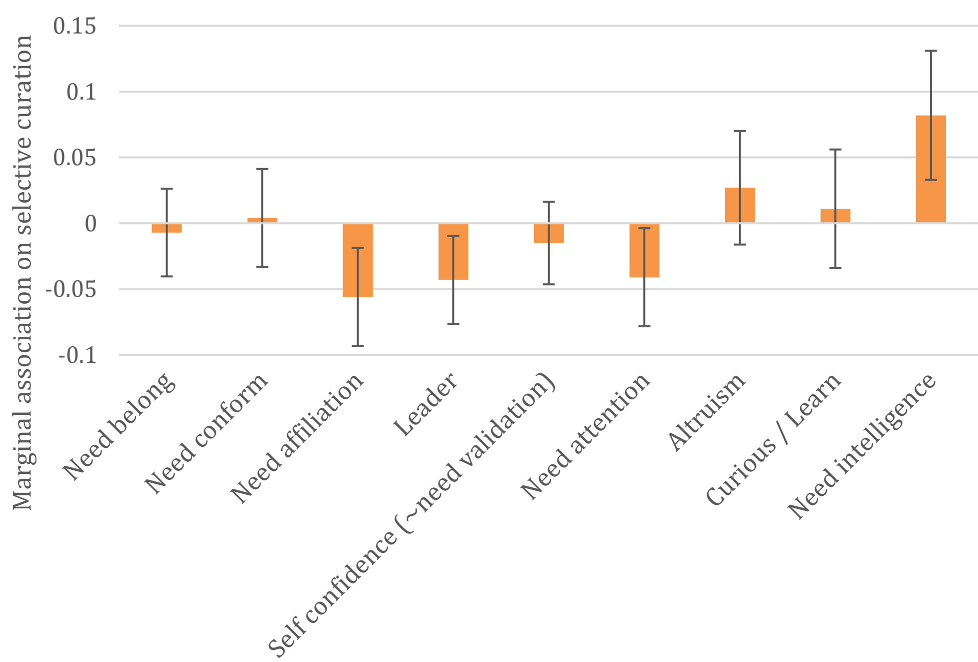

Marginal association of traits on subjects' propensity to share party-favorable stories (selective curation).

Bibliography

@mjahr. "Never miss important Tweets from people you follow." Blog.twitter.com (blog). February 10, 2016. https://blog.twitter.com/official/en_us/a/2016/never-miss-important-tweets-from-people-you-follow.html.

Abrams v. United States, 250 U.S. 616 (1919).

Acquisti, Alessandro, and Ralph Gross. "Imagined communities: Awareness, information sharing, and privacy on the Facebook." In Privacy Enhancing Technologies. PET 2006. Lecture Notes in Computer Science, vol 4258. edited by G. Danezis and P. Golle, 36-58. Berlin, Heidelberg: Springer, 2006.

Adamic, Lada A., and Natalie Glance. "The political blogosphere and the 2004 US election: Divided they blog." In *Proceedings of the 3rd International Workshop on Link Discovery*, 36-43. ACM, 2005.

Allcott, Hunt, and Matthew Gentzkow. "Social media and fake news in the 2016 election." *Journal of Economic Perspectives* 31, no. 2 (2017): 211-36.

An, Jisun, Daniele Quercia, Meeyoung Cha, Krishna Gummadi, and Jon Crowcroft. "Sharing political news: The balancing act of intimacy and socialization in selective exposure." *EPJ Data Science* 3, no. 1 (2014a): 12.

An, Jisun, Daniele Quercia, and Jon Crowcroft. "Partisan sharing: Facebook evidence and societal consequences." In *Proceedings of the second ACM Conference on Online Social Networks*, 13-24. ACM, 2014b.

Anderson, Ashley A., Dominique Brossard, Dietram A. Scheufele, Michael A. Xenos, and Peter Ladwig. "The “nasty effect:” Online incivility and risk perceptions of emerging technologies." *Journal of Computer-Mediated Communication* 19, no. 3 (2014): 373-387.

Back, Mitja D., Juliane M. Stopfer, Simine Vazire, Sam Gaddis, Stefan C. Schmukle, Boris Egloff, and Samuel D. Gosling. "Facebook profiles reflect actual personality, not self-idealization." *Psychological Science* 21, no. 3 (2010): 372-374.

Bail, Christopher A., Lisa Argyle, Taylor Brown, John Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. "Exposure to opposing views can increase political polarization: Evidence from a large-scale field experiment on social media." Proceedings of the National Academy of Sciences 37, no. 115 (2018): 9216-9221.

Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. "Exposure to ideologically diverse news and opinion on Facebook." *Science* 348, no. 6239 (2015): 1130-1132.

Baldassarri, Delia, and Peter Bearman. "Dynamics of political polarization." *American Sociological Review* 72, no. 5 (2007): 784-811.

Banisch, Sven, Ricardo Lima, and Tanya Araújo. "Agent based models and opinion dynamics as Markov chains." *Social Networks* 34, no. 4 (2012): 549-561.

Baños, Raquel, Javier Borge-Holthoefer, Ning Wang, Yamir Moreno, and Sandra González-Bailón. "Diffusion dynamics with changing network composition." *Entropy* 15, no. 11 (2013): 4553-4568.

Barabas, Jason, and Jennifer Jerit. "Estimating the causal effects of media coverage on policy-specific knowledge." *American Journal of Political Science* 53, no. 1 (2009): 73-89.

Barasch, Alixandra, and Jonah Berger. "Broadcasting and narrowcasting: How audience size affects what people share." *Journal of Marketing Research* 51, no. 3 (2014): 286-299.

Barberá, Pablo. "How social media reduces mass political polarization. Evidence from Germany, Spain, and the US." *Job Market Paper, New York University* (2014).

Barberá, Pablo. "Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data." *Political Analysis* 23, no. 1 (2015): 76-91.

Barberá, Pablo, John T. Jost, Jonathan Nagler, Joshua A. Tucker, and Richard Bonneau. "Tweeting from left to right: Is online political communication more than an echo chamber?" *Psychological Science* 26, no. 10 (2015): 1531-1542.

Bargh, John A., and Tanya L. Chartrand. "The unbearable automaticity of being." *American Psychologist* 54, no. 7 (1999): 462-479.

Baum, Matthew A., and Tim Groeling. "New media and the polarization of American political discourse." *Political Communication* 25, no. 4 (2008): 345-365.

Baumeister, Roy F., and E. J. Masicampo. "Conscious thought is for facilitating social and cultural interactions: How mental simulations serve the animal–culture interface." *Psychological Review* 117, no. 3 (2010): 945.

Baumeister, Roy F., and Mark R. Leary. "The need to belong: desire for interpersonal attachments as a fundamental human motivation." *Psychological Bulletin* 117, no. 3 (1995): 497-529.

Bazarova, Natalya N., and Yoon Hyung Choi. "Self-disclosure in social media: Extending the functional approach to disclosure motivations and characteristics on social network sites." *Journal of Communication* 64, no. 4 (2014): 635-657.

Bednar, Jenna, Aaron Bramson, Andrea Jones-Rooy, and Scott Page. "Emergent cultural signatures and persistent diversity: A model of conformity and consistency." *Rationality and Society* 22, no. 4 (2010): 407-444.

Ben-Nun Bloom, Pazit, and Lindsey Clark Levitan. "We're closer than I thought: Social network heterogeneity, morality, and political persuasion." *Political Psychology* 32, no. 4 (2011): 643-665.

Benkler, Yochai. The Wealth of Networks: How Social Production Transforms Markets and Freedom. New Haven: Yale University Press, 2006.

Bennett, W. Lance, and Shanto Iyengar. "A new era of minimal effects? The changing foundations of political communication." *Journal of Communication* 58, no. 4 (2008): 707-731.

Berger, Jonah. "Word of mouth and interpersonal communication: A review and directions for future research." *Journal of Consumer Psychology* 24, no. 4 (2014): 586-607.

Bimber, Bruce. "The Internet and political transformation: Populism, community, and accelerated pluralism." *Polity* 31, no. 1 (1998): 133-160.

Moore, C.R. and M.A. Crotty. "Australian Masculinities." In *International Encyclopedia of Men and Masculinities*, edited by M. Flood, J.K. Gardiner, B. Pease, and K. Pringle. Routledge, 2007.

Bor, Stephanie E. "Echo Chamber Phenomenon." In *Encyclopedia of Social Media and Politics*, edited by K. Harvey. Sage Publishing, 2014

Boutyline, Andrei, and Robb Willer. "The social structure of political echo chambers: Variation in ideological homophily in online networks." *Political Psychology* 38, no. 3 (2017): 551-569.

Boyd, Danah M., and Nicole B. Ellison. "Social network sites: Definition, history, and scholarship." *Journal of Computer‐Mediated Communication* 13, no. 1 (2007): 210-230.

Brady, William J., Julian A. Wills, John T. Jost, Joshua A. Tucker, and Jay J. Van Bavel. "Emotion shapes the diffusion of moralized content in social networks." *Proceedings of the National Academy of Sciences* 114, no. 28 (2017): 7313-7318.

Broh, C. Anthony. "Horse-race journalism: Reporting the polls in the 1976 presidential election." *Public Opinion Quarterly* 44, no. 4 (1980): 514-529.

Brown, Jacqueline Johnson, and Peter H. Reingen. "Social ties and word-of-mouth referral behavior." *Journal of Consumer Research* 14, no. 3 (1987): 350-362.

Brownlee, Chip. "YouTube's New Crackdown on Hate Speech Comes at a Curious Time." Slate, June 5, 2019. https://slate.com/technology/2019/06/youtube-hate-speech-crackdown-steven-crowder-carlos-maza.html

Brundidge, Jennifer. "Encountering “difference” in the contemporary public sphere: The contribution of the Internet to the heterogeneity of political discussion networks." *Journal of Communication* 60, no. 4 (2010): 680-700.

Bullock, John G., Alan S. Gerber, Seth J. Hill, and Gregory A. Huber. *Partisan bias in factual beliefs about politics*. No. w19080. National Bureau of Economic Research, 2013.

Cameron, James E. "A three-factor model of social identity." *Self and Identity* 3, no. 3 (2004): 239-262.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. *The American Voter*. University of Chicago Press, 1960.

Cappella, Joseph N., Hyun Suk Kim, and Dolores Albarracín. "Selection and transmission processes for information in the emerging media environment: Psychological motives and message characteristics." *Media Psychology* 18, no. 3 (2015): 396-424.

Centola, Damon, and Michael Macy. "Complex contagions and the weakness of long ties." *American Journal of Sociology* 113, no. 3 (2007): 702-734.

Chiu, Chao-Min, Meng-Hsiang Hsu, and Eric TG Wang. "Understanding knowledge sharing in virtual communities: An integration of social capital and social cognitive theories." *Decision Support Systems* 42, no. 3 (2006): 1872-1888.

Clifford, Scott, Ryan M. Jewell, and Philip D. Waggoner. "Are samples drawn from Mechanical Turk valid for research on political ideology?" *Research & Politics* 2, no. 4 (2015): 1-9.

Cohen, Geoffrey L., David K. Sherman, Anthony Bastardi, Lillian Hsu, Michelle McGoey, and Lee Ross. "Bridging the partisan divide: Self-affirmation reduces ideological closed-mindedness and inflexibility in negotiation." *Journal of Personality and Social Psychology* 93, no. 3 (2007): 415-430.

Colleoni, Elanor, Alessandro Rozza, and Adam Arvidsson. "Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data." *Journal of Communication* 64, no. 2 (2014): 317-332.

Conover, Pamela Johnston, and Stanley Feldman. "The origins and meaning of liberal/conservative self-identifications." *American Journal of Political Science* 25, no. 4 (1981): 617-645.

Conover, Michael D., Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. "Political polarization on twitter." In *Fifth International AAAI Conference on Weblogs and Social Media*. 2011.

Constine, Josh. "Facebook changes feed to promote posts that aren't fake, sensational, or spam." Techcrunch, January 31, 2017. https://techcrunch.com/2017/01/31/facebook-authentic-news/

Crockett, M. J. "Moral outrage in the digital age." *Nature Human Behaviour* 1, no. 11 (2017): 769-771.

Damasio, Antonio R. Descartes' Error: Emotion, Reason, and the Human Brain. New York: Putnam, 1994.

Dandekar, Pranav, Ashish Goel, and David T. Lee. "Biased assimilation, homophily, and the dynamics of polarization." *Proceedings of the National Academy of Sciences* 110, no. 15 (2013): 5791-5796.

Dunbar, Robin IM. "The social brain hypothesis." Evolutionary Anthropology: Issues, News, and Reviews: Issues, News, and Reviews 6, no. 5 (1998): 178-190.

Earl, Allison, William Hart, Alex Burton, and Dolores Albarracin, D. "Biased Information Selection on Behalf of Others: The Case for Vicarious Selective Exposure." Working paper. 2018.

Eisenberger, Naomi I., Matthew D. Lieberman, and Kipling D. Williams. "Does rejection hurt? An fMRI study of social exclusion." *Science* 302, no. 5643 (2003): 290-292.

Ellison, Nicole B., and Jessica Vitak. "Social network site affordances and their relationship to social capital processes." *The Handbook of the Psychology of Communication Technology* 32 (2015): 205-228.

Ellsworth, Phoebe C., and Klaus R. Scherer. "Appraisal processes in emotion." *Handbook of Affective Sciences* (2003): 572-595.

Emba, Christine. "Confirmed: Echo chambers exist on social media. So what do we do about them?" Washington Post, July 14, 2016. https://www.washingtonpost.com/news/in-theory/wp/2016/07/14/confirmed-echo-chambers-exist-on-social-media-but-what-can-we-do-about-them/

Erdős, Paul, and Alfréd Rényi. "On the evolution of random graphs." *Publication of the Mathematical Institute of the Hungarian Academy of Sciences* 5, no. 1 (1960): 17-60.

Erisen, Elif, and Cengiz Erisen. "The effect of social networks on the quality of political thinking." *Political Psychology* 33, no. 6 (2012): 839-865.

Faris, R., H. Roberts, B. Etling, N. Bourassa, E. Zuckerman, and Y. Benkler. "Partisanship, propaganda and disinformation: Online media and the 2016 US presidential election." Cambridge: Berkman Klein Center, 2017.

Farrell, Henry. "The Consequences of the Internet for Politics." *Annual Review of Political Science* 15, no. 1 (2012), 35–52.

Festinger, Leon. "Cognitive dissonance." *Scientific American* 207, no. 4 (1962): 93-106.

Festinger, Leon, and James M. Carlsmith. "Cognitive consequences of forced compliance." *The Journal of Abnormal and Social Psychology* 58, no. 2 (1959): 203-210.

Fiorina, Morris P., Samuel J. Abrams, and Jeremy C. Pope. Culture war? *The myth of a polarized America.* New York: Pearson Longman, 2005.

Jory, Patrick. "Republicanism in Thai History." In *A Sarong for Clio: Essays on the Intellectual and Cultural History of Thailand*, edited by Maurizio Peleggi, 97-117. Ithaca, NY: Cornell University Press, 2015.

Fiorina, Morris P. "A dark side of civic engagement." In *Civic Engagement in American Democracy,* edited by Theda Skocpol and Morris Fiorina, 319-352. Washington D.C: Brookings Institute Press, 1999.

Flache, Andreas, and Michael W. Macy. "Local convergence and global diversity: From interpersonal to social influence." *Journal of Conflict Resolution* 55, no. 6 (2011): 970-995.

Flaxman, Seth, Sharad Goel, and Justin M. Rao. "Filter bubbles, echo chambers, and online news consumption." *Public Opinion Quarterly* 80, no. 1 (2016): 298-320.

Frayer, Lauren. "How The Spread Of Fake Stories In India Has Led To Violence." NPR. July 17, 2018. https://www.npr.org/2018/07/17/629896525/how-the-spread-of-fake-stories-in-india-has-led-to-violence

Frey, Dieter. "Recent research on selective exposure to information." *Advances in Experimental Social Psychology* 19 (1986): 41-80.

Gans, Herbert J. *Deciding what's news*. New York: Vintage Books, 1979.

Garrett, R. Kelly. "Echo chambers online?: Politically motivated selective exposure among Internet news users." *Journal of Computer-Mediated Communication* 14, no. 2 (2009): 265-285.

Garrett, R. Kelly, Shira Dvir Gvirsman, Benjamin K. Johnson, Yariv Tsfati, Rachel Neo, and Aysenur Dal. "Implications of pro-and counterattitudinal information exposure for affective polarization." *Human Communication Research* 40, no. 3 (2014): 309-332.

Gaver, William W. "Technology affordances." In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 79-84. ACM Press, 1991.

Gentzkow, Matthew, and Jesse M. Shapiro. "Ideological segregation online and offline." *The Quarterly Journal of Economics* 126, no. 4 (2011): 1799-1839.

Gibson, James J. "The theory of affordances." In *Perceiving, Acting, and Knowing*, edited by R.E. Shaw and J. Bransford. Hilldale: Lawrence Erlbaum Assoc., 1977.

Goel, Sharad, Winter Mason, and Duncan J. Watts. "Real and perceived attitude agreement in social networks." *Journal of Personality and Social Psychology* 99, no. 4 (2010): 611-621.

Goffman, Erving. The Presentation of Self in Everyday Life. Doubleday, 1959.

Gollwitzer, Peter M. "Implementation intentions: strong effects of simple plans." *American Psychologist* 54, no. 7 (1999): 493-503.

Goldenberg, Jacob, Barak Libai, and Eitan Muller. "Talk of the network: A complex systems look at the underlying process of word-of-mouth." *Marketing Letters* 12, no. 3 (2001): 211-223.

Google Trends. https://trends.google.com/trends/explore?date=2016-11-01%202016-11-15&geo=US&q=liberal%20bubble

Gosling, Samuel D., Sam Gaddis, and Simine Vazire. "Personality impressions based on facebook profiles." *International Conference on Web and Social Media* 7, (2007): 1-4.

Grabowicz, Przemyslaw A., José J. Ramasco, Esteban Moro, Josep M. Pujol, and Victor M. Eguiluz. "Social features of online networks: The strength of intermediary ties in online social media." *PloS One* 7, no. 1 (2012): e29358.

Green, Donald P., Bradley Palmquist, and Eric Schickler. *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. Yale University Press, 2004.

Guess, Andrew M. "Measure for measure: An experimental test of online political media exposure." *Political Analysis* 23, no. 1 (2015): 59-75.

Guess, Andrew, Brendan Nyhan, and Jason Reifler. "Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign." *European Research Council* 9 (2018).

Haidt, Jonathan. Interview with Sean Illing. "Why Social Media is Terrible for Multiethnic Societies." Vox, Updated January 5, 2018. https://www.vox.com/policy-and-politics/2016/11/15/13593670/donald-trump-social-media-culture-politics

Halberstam, Yosh, and Brian Knight. "Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter." *Journal of Public Economics* 143 (2016): 73-88.

Hanna, Alexander, Ben Sayre, Leticia Bode, JungHwan Yang, and Dhavan Shah. "Mapping the political Twitterverse: Candidates and their followers in the midterms." In *Fifth International AAAI Conference on Weblogs and Social Media*. 2011.

Hart, William, Dolores Albarracín, Alice H. Eagly, Inge Brechan, Matthew J. Lindberg, and Lisa Merrill. "Feeling validated versus being correct: a meta-analysis of selective exposure to information." *Psychological Bulletin* 135, no. 4 (2009): 555-588.

Herrman, John. "Inside Facebook's (Totally Insane, Unintentionally Gigantic, Hyperpartisan) Political-Media Machine." New York Times. August 24, 2016. https://www.nytimes.com/2016/08/28/magazine/inside-facebooks-totally-insane-unintentionally-gigantic-hyperpartisan-political-media-machine.html

Hetherington, Marc J. "Putting polarization in perspective." *British Journal of Political Science* 39, no. 2 (2009): 413-448.

Hibbing, John R., and Elizabeth Theiss-Morse. *Stealth Democracy: Americans' Beliefs about How Government Should Work*. Cambridge University Press, 2002.

Hill, Craig A. "Affiliation motivation: people who need people… but in different ways." *Journal of personality and social psychology* 52, no. 5 (1987): 1008-1018.

Himelboim, Itai, Stephen McCreery, and Marc Smith. "Birds of a feather tweet together: Integrating network and content analyses to examine cross-ideology exposure on Twitter." *Journal of Computer-Mediated Communication* 18, no. 2 (2013): 154-174.

Hogg, Michael A., and Scott A. Reid. "Social identity, self-categorization, and the communication of group norms." *Communication Theory* 16, no. 1 (2006): 7-30.

Hogg, Michael A., Deborah J. Terry, and Katherine M. White. "A tale of two theories: A critical comparison of identity theory with social identity theory." *Social Psychology Quarterly* (1995): 255-269.

Holbrook, Allyson L., and Jon A. Krosnick. "Social desirability bias in voter turnout reports: Tests using the item count technique." *Public Opinion Quarterly* 74, no. 1 (2009): 37-67.

Hosanakar, Kartik. "Blame the Echo Chamber on Facebook. But Blame Yourself, Too." Wired. November 24, 2016. https://www.wired.com/2016/11/facebook-echo-chamber/

Hsu, Meng-Hsiang, Teresa L. Ju, Chia-Hui Yen, and Chun-Ming Chang. "Knowledge sharing behavior in virtual communities: The relationship between trust, self-efficacy, and outcome expectations." *International Journal of Human-Computer Studies* 65, no. 2 (2007): 153-169.

Huber, Gregory Alain, and Neil Ankur Malhotra. Dimensions of political homophily: Isolating choice homophily along political characteristics. Stanford Graduate School of Business, 2013.

Huckfeldt, Robert, and John Sprague. "Networks in context: The social flow of political information." *American Political Science Review* 81, no. 4 (1987): 1197-1216.

Huckfeldt, Robert, Paul Allen Beck, Russell J. Dalton, and Jeffrey Levine. "Political environments, cohesive social groups, and the communication of public opinion." *American Journal of Political Science* 39, no. 4 (1995): 1025-1054.

Huckfeldt, Robert, Jeanette Morehouse Mendez, and Tracy Osborn. "Disagreement, ambivalence, and engagement: The political consequences of heterogeneous networks." *Political Psychology* 25, no. 1 (2004): 65-95.

Huddy, Leonie. "From social to political identity: A critical examination of social identity theory." *Political Psychology* 22, no. 1 (2001): 127-156.

Ingram, Matthew. "Facebook has taken over from Google as a traffic source for news." Fortune. August 18, 2015. http://fortune.com/2015/08/18/facebook-google/

Iyengar, Shanto, and Kyu S. Hahn. "Red media, blue media: Evidence of ideological selectivity in media use." *Journal of Communication* 59, no. 1 (2009): 19-39.

Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. "Affect, not ideology a social identity perspective on polarization." *Public Opinion Quarterly* 76, no. 3 (2012): 405-431.

Jackson, Melinda S. "Priming the sleeping giant: The dynamics of Latino political identity and vote choice." *Political Psychology* 32, no. 4 (2011): 691-716.

Jungherr, Andreas. "Twitter in politics: a comprehensive literature review." *Available at SSRN 2865150* (2014).

Kankanhalli, Atreyi, Bernard CY Tan, and Kwok-Kee Wei. "Contributing knowledge to electronic knowledge repositories: An empirical investigation." *MIS Quarterly* 29, no. 1 (2005): 133-143.

Katz, Elihu, and Paul F. Lazarsfeld. *Personal Influence.* New York: Free Press, 1955.

Keltner, Dacher, and James J. Gross. "Functional accounts of emotions." *Cognition & Emotion* 13, no. 5 (1999): 467-480.

Klar, Samara. "The influence of competing identity primes on political preferences." *The Journal of Politics* 75, no. 4 (2013): 1108-1124.

Klar, Samara, and Yanna Krupnikov. *Independent Politics*. Cambridge University Press, 2016.

Klayman, Joshua. "Varieties of confirmation bias." In *Psychology of Learning and Motivation*, vol. 32, pp. 385-418. Academic Press, 1995.

Knowles, Megan L., and Wendi L. Gardner. "Benefits of membership: The activation and amplification of group identities in response to social rejection." *Personality and Social Psychology Bulletin* 34, no. 9 (2008): 1200-1213.

Kunda, Ziva. "The case for motivated reasoning." *Psychological Bulletin* 108, no. 3 (1990): 480-498.

Kurzban, Robert. Why everyone (else) is a hypocrite: Evolution and the modular mind. Princeton University Press, 2012.

Kurzban, Robert, and C. Athena Aktipis. "Modularity and the social mind: Are psychologists too self-ish?." *Personality and Social Psychology Review* 11, no. 2 (2007): 131-149.

Lane, Robert Edwards. *Political Ideology: Why the American Common Man Believes What he Does*. Oxford, England: Free Press of Glencoe, 1962.

Lau, Richard R., and David P. Redlawsk. "Voting correctly." *American Political Science Review* 91, no. 3 (1997): 585-598.

Lazarsfeld, Paul Felix, Bernard Berelson, and Hazel Gaudet. *The People's Choice*. Oxford, England: Duell, Sloan & Pearce, 1944.

Lazer, David, Brian Rubineau, Carol Chetkovich, Nancy Katz, and Michael Neblo. "The coevolution of networks and political attitudes." *Political Communication* 27, no. 3 (2010): 248-274.

Lazer, David MJ, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, and Miriam J. Metzger. "The science of fake news." *Science* 359, no. 6380 (2018): 1094-1096.

Leary, Mark R., Kristine M. Kelly, Catherine A. Cottrell, and Lisa S. Schreindorfer. "Construct validity of the need to belong scale: Mapping the nomological network." *Journal of Personality Assessment* 95, no. 6 (2013): 610-624.

Leary, Mark R., Katrina P. Jongman-Sereno, and Kate J. Diebels. "Measures of concerns with public image and social evaluation." In *Measures of Personality and Social Psychological constructs*, pp. 448-473. Academic Press, 2015.

Lelkes, Yphtach, Gaurav Sood, and Shanto Iyengar. "The hostile audience: The effect of access to broadband internet on partisan affect." *American Journal of Political Science* 61, no. 1 (2017): 5-20.

Levay, Kevin E., Jeremy Freese, and James N. Druckman. "The demographic and political composition of Mechanical Turk samples." *Sage Open* 6, no. 1 (2016): 1-17.

Levendusky, Matthew. "Partisan media exposure and attitudes toward the opposition." *Political Communication* 30, no. 4 (2013): 565-581.

Levendusky, Matthew. "Americans, not partisans: Can priming American national identity reduce affective polarization?." *The Journal of Politics* 80, no. 1 (2018): 59-70.

Lieberman, Matthew D. Social: Why our brains are wired to connect. OUP Oxford, 2013.

Lin, Hsiu-Fen. "Effects of extrinsic and intrinsic motivation on employee knowledge sharing intentions." *Journal of Information Science* 33, no. 2 (2007): 135-149.

Lippman, Walter. *Public Opinion*. New York: Macmillan, 1922.

Litt, Eden. "Knock, knock. Who's there? The imagined audience." *Journal of Broadcasting & Electronic Media* 56, no. 3 (2012): 330-345.

Lord, Charles G., Lee Ross, and Mark R. Lepper. "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence." *Journal of Personality and Social Psychology* 37, no. 11 (1979): 2098-2109.

Lum, Nick. "Tools for Bursting Your Filter Bubble." Medium. January 31, 2017. https://medium.com/@nicklum/tools-for-bursting-your-filter-bubble-51f482027128

Lupia, Arthur. "Shortcuts versus encyclopedias: Information and voting behavior in California insurance reform elections." *American Political Science Review* 88, no. 1 (1994): 63-76.

Maccoby, Eleanor E., and Nathan Maccoby. "The interview: A tool of social science." *Handbook of Social Psychology* 1, no. 1 (1954): 449-487.

Malka, Ariel, and Yphtach Lelkes. "More than ideology: Conservative–liberal identity and receptivity to political cues." *Social Justice Research* 23, no. 2-3 (2010): 156-188.

Manin, Bernard. "On legitimacy and political deliberation." *Political Theory* 15, no. 3 (1987): 338-368.

Manjoo, Farhad. True Enough: Learning to Live in a Post-fact Society. John Wiley & Sons, 2011.

Marcus, George E. "Emotions in politics." *Annual Review of Political Science* 3, no. 1 (2000): 221-250.

Marwick, Alice E., and Danah Boyd. "I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience." *New Media & Society* 13, no. 1 (2011): 114-133.

Mason, Lilliana, and Julie Wronski. "One tribe to bind them all: How our social group attachments strengthen partisanship." *Political Psychology* 39 (2018): 257-277.

McClain, Paula D., Jessica D. Johnson Carew, Eugene Walton Jr, and Candis S. Watts. "Group membership, group identity, and group consciousness: Measures of racial identity in American politics?" *Annual Review of Political Science* 12 (2009): 471-485.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. "Birds of a feather: Homophily in social networks." *Annual Review of Sociology* 27, no. 1 (2001): 415-444.

Mendelberg, Tali. "The deliberative citizen: Theory and evidence." *Political Decision Making, Deliberation and Participation* 6, no. 1 (2002): 151-193.

Mill, John Stuart. *On Liberty.* Longmans, Green, Reader, and Dyer, 1869.

Millham, Jim, and Richard W. Kellogg. "Need for social approval: impression management or self-deception?." *Journal of Research in Personality* 14, no. 4 (1980): 445-457.

Mønsted, Bjarke, Piotr Sapieżyński, Emilio Ferrara, and Sune Lehmann. "Evidence of complex contagion of information in social media: An experiment using Twitter bots." *PloS One* 12, no. 9 (2017): e0184148.

Moore, Michael. Facebook (post). November 9, 2016. https://www.facebook.com/mmflint/posts/10153913074756857

Mosseri, Adam. "News Feed FYI: Bringing People Closer Together." Facebook Business (blog). January 11, 2018. https://www.facebook.com/business/news/news-feed-fyi-bringing-people-closer-together

Mutz, Diana C. "Cross-cutting social networks: Testing democratic theory in practice." *American Political Science Review* 96, no. 1 (2002a): 111-126.

Mutz, Diana C. "The consequences of cross-cutting networks for political participation." *American Journal of Political Science* (2002b): 838-855.

Mutz, Diana C., and Jeffery J. Mondak. "The workplace as a context for cross-cutting political discourse." *The Journal of Politics* 68, no. 1 (2006): 140-155.

Nadkarni, Ashwini, and Stefan G. Hofmann. "Why do people use Facebook?." *Personality and Individual Differences* 52, no. 3 (2012): 243-249.

Negroponte, Nicholas. *Being digital*. Vintage, 1996.

Nelson, Thomas E., Gregory Gwiasda, and Joseph Lyons. "Vilification and values." *Political Psychology* 32, no. 5 (2011): 813-835.

Neuberg, Steven L., and Catherine A. Cottrell. "Managing the threats and opportunities afforded by human sociality." *Group Dynamics: Theory, Research, and Practice* 12, no. 1 (2008): 63-72.

Nichols, Austin Lee, and Gregory D. Webster. "The single-item need to belong scale." *Personality and Individual Differences* 55, no. 2 (2013): 189-192.

Nickerson, David W., and Todd Rogers. "Do you have a voting plan? Implementation intentions, voter turnout, and organic plan making." *Psychological Science* 21, no. 2 (2010): 194-199.

Nisbett, Richard E., and Timothy D. Wilson. "Telling more than we can know: Verbal reports on mental processes." *Psychological Review* 84, no. 3 (1977): 231-259.

NPR Staff. "The Reason Your Feed Became An Echo Chamber — And What To Do About It." NPR, July 24, 2016. https://www.npr.org/sections/alltechconsidered/2016/07/24/486941582/the-reason-your-feed-became-an-echo-chamber-and-what-to-do-about-it

Noel, Jeffrey G., Daniel L. Wann, and Nyla R. Branscombe. "Peripheral ingroup membership status and public negativity toward outgroups." *Journal of personality and social psychology* 68, no. 1 (1995): 127-137.

Nyhan, Brendan, and Jason Reifler. "When corrections fail: The persistence of political misperceptions." *Political Behavior* 32, no. 2 (2010): 303-330.

Orbach, Maya, Maegen Demko, Jeremy Doyle, Benjamin N. Waber, and Alex Pentland. "Sensing informal networks in organizations." *American Behavioral Scientist* 59, no. 4 (2015): 508-524.

Orne, Martin T. "On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications." *American Psychologist* 17, no. 11 (1962): 776-783.

Page, Scott E. The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies-New Edition. Princeton University Press, 2008.

Pariser, Eli. The Filter Bubble: How the New Personalized Web is Changing What we Read and How we Think. Penguin, 2011.

Parsons, Bryan M. "Social networks and the affective impact of political disagreement." *Political Behavior* 32, no. 2 (2010): 181-204.

Piore, Adam. "Technologists are trying to fix the “filter bubble” problem that tech helped create." Technologyreview.com, August 22, 2018. https://www.technologyreview.com/s/611826/technologists-are-trying-to-fix-the-filter-bubble-problem-that-tech-helped-create/

Prior, Markus. "News vs. entertainment: How increasing media choice widens gaps in political knowledge and turnout." *American Journal of Political Science* 49, no. 3 (2005): 577-592.

Prior, Markus. "Media and political polarization." *Annual Review of Political Science* 16 (2013): 101-127.

Pronin, Emily, Daniel Y. Lin, and Lee Ross. "The bias blind spot: Perceptions of bias in self versus others." *Personality and Social Psychology Bulletin* 28, no. 3 (2002): 369-381.

Redlawsk, David P., Andrew JW Civettini, and Karen M. Emmerson. "The affective tipping point: Do motivated reasoners ever “get it”?." *Political Psychology* 31, no. 4 (2010): 563-593.

Remnick, David. "Obama Reckons with a Trump Presidency." New Yorker, November 28, 2016. https://www.newyorker.com/magazine/2016/11/28/obama-reckons-with-a-trump-presidency

Riker, William H., and Peter C. Ordeshook. "A Theory of the Calculus of Voting." *American Political Science Review* 62, no. 1 (1968): 25-42.

Rimé, Bernard. "Emotion elicits the social sharing of emotion: Theory and empirical review." *Emotion Review* 1, no. 1 (2009): 60-85.

Romero, Daniel M., Brendan Meeder, and Jon Kleinberg. "Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter." In *Proceedings of the 20th international conference on World wide web*, pp. 695-704. ACM, 2011.

Romm, Tony & Elizabeth Dwoskin. "CEO says he's rethinking the core of how Twitter works." Washington Post, August 15, 2018. https://www.washingtonpost.com/technology/2018/08/15/jack-dorsey-says-hes-rethinking-core-how-twitter-works/

Romm, Tony & Craig Timberg. "Facebook and Twitter testified before Congress. Conservative conspiracy theorists lurked behind them." September 5, 2018. https://www.washingtonpost.com/technology/2018/09/05/facebook-twitter-sandberg-dorsey-congress-tech-hearings/

Ross, Lee. "The intuitive psychologist and his shortcomings: Distortions in the attribution process." In *Advances in Experimental Social Psychology*, vol. 10, pp. 173-220. Academic Press, 1977.

Sasahara, Kazutoshi, Wen Chen, Hao Peng, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. "On the Inevitability of Online Echo Chambers." *ArXiv*, (2019).

Schelling, Thomas C. "Dynamic models of segregation." *Journal of Mathematical Sociology* 1, no. 2 (1971): 143-186.

Schudson, Michael. Discovering the News: A Social History of American Newspapers. Basic Books, 1981.

Schwarz, Norbert. "Self-reports: how the questions shape the answers." *American Psychologist* 54, no. 2 (1999): 93-105.

Sears, David O., and Jonathan L. Freedman. "Selective exposure to information: A critical review." *Public Opinion Quarterly* 31, no. 2 (1967): 194-213.

Settle, Jaime. ’Frenemies: How Social Media Polarizes America.’ Presented at APSA Political Psychology Preconference. Philadelphia, 2016.

Settle, Jaime E. *Frenemies: How Social Media Polarizes America*. Cambridge University Press, 2018.

Shin, Jieun, and Kjerstin Thorson. "Partisan selective sharing: The biased diffusion of fact-checking messages on social media." *Journal of Communication* 67, no. 2 (2017): 233-255.

Shearer, Elisa, and Jeffrey Gottfried. "News use across social media platforms 2017." *Pew Research Center* 7 (2017).

Shoemaker, Pamela J., Martin Eichholz, Eunyi Kim, and Brenda Wrigley. "Individual and routine forces in gatekeeping." *Journalism & Mass Communication Quarterly* 78, no. 2 (2001): 233-246.

Siegel, David A. "Social networks and the mass media." *American Political Science Review* 107, no. 4 (2013): 786-805.

Sigal, Leon V. Reporters and Officials: The Organization and Politics of Newsmaking. DC Heath, 1973.

Silverman, Craig. "This Analysis Shows How Viral Fake Election News Stories Outperformed Real News On Facebook." Buzzfeed, November 16, 2016. https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook

Saturday Night Live. "The Bubble." Youtube video. November 20, 2016. https://www.youtube.com/watch?v=vKOb-kmOgpI

Sobieraj, Sarah, Jeffrey M. Berry, and Amy Connors. "Outrageous political opinion and political anxiety in the US." *Poetics* 41, no. 5 (2013): 407-432.

Solon, Olivia & Sabrina Siddiqui."Russia-backed Facebook posts 'reached 126m Americans' during US election." The Guardian, October 30, 2017. https://www.theguardian.com/technology/2017/oct/30/facebook-russia-fake-accounts-126-million

Soroka, Stuart N. "The gatekeeping function: Distributions of information in media and the real world." *The Journal of Politics* 74, no. 2 (2012): 514-528.

Soroka, Stuart, Mark Daku, Dan Hiaeshutter-Rice, Lauren Guggenheim, and Josh Pasek. "Negativity and positivity biases in economic news coverage: Traditional versus social media." *Communication Research* 45, no. 7 (2018): 1078-1098.

Steele, Claude M. "The psychology of self-affirmation: Sustaining the integrity of the self." In *Advances in Experimental Social Psychology*, vol. 21, pp. 261-302. Academic Press, 1988.

Stieglitz, Stefan, and Linh Dang-Xuan. "Political communication and influence through microblogging--An empirical analysis of sentiment in Twitter messages and retweet behavior." In *2012 45th Hawaii International Conference on System Sciences*, pp. 3500-3509. IEEE, 2012.

Skitka, Linda J. "Do the means always justify the ends, or do the ends sometimes justify the means? A value protection model of justice reasoning." *Personality and Social Psychology Bulletin* 28, no. 5 (2002): 588-597.

Stroud, Natalie Jomini. "Media use and political predispositions: Revisiting the concept of selective exposure." *Political Behavior* 30, no. 3 (2008): 341-366.

Suhay, Elizabeth. "Explaining group influence: The role of identity and emotion in political conformity and polarization." *Political Behavior* 37, no. 1 (2015): 221-251.

Sunstein, Cass R. *Republic.com 2.0*. Princeton University Press, 2009.

Sunstein, Cass R. # Republic: Divided Democracy in the Age of Social Media. Princeton University Press, 2018.

Swasy, Alecia. How Journalists Use Twitter: The Changing Landscape of US Newsrooms. Lexington Books, 2016.

Taber, Charles S., and Milton Lodge. "Motivated skepticism in the evaluation of political beliefs." *American Journal of Political Science* 50, no. 3 (2006): 755-769.

Tajfel, Henri. "Social identity and intergroup behaviour." *Information (International Social Science Council)* 13, no. 2 (1974): 65-93.

Taub, Amanda & Max Fisher. "Where Countries Are Tinderboxes and Facebook Is a Match." New York Times, April 21. 2018. https://www.nytimes.com/2018/04/21/world/asia/facebook-sri-lanka-riots.html

Taylor, J. Benjamin. "The Educative Effects of Extreme Television Media." In *Extreme Media and American Politics*, pp. 55-92. Palgrave Macmillan, Cham, 2017.

Tetlock, Philip E., and Antony S. Manstead. "Impression management versus intrapsychic explanations in social psychology: A useful dichotomy?." *Psychological Review* 92, no. 1 (1985): 59-77.

Thaler, Richard H., and Cass R. Sunstein. Nudge: Improving Decisions about Health, Wealth, and Happiness. Penguin, 2009.

The Economist staff. "Google rewards reputable reporting, not left-wing politics." The Economist, June 8, 2019. https://www.economist.com/graphic-detail/2019/06/08/google-rewards-reputable-reporting-not-left-wing-politics

Tingley, Dustin, Teppei Yamamoto, Kentaro Hirose, Luke Keele, and Kosuke Imai. "Mediation: R Package for Causal Mediation Analysis." *Journal of Statistical Software* 59, no. 5 (2014): 1-38.

Tooby, John, and Leda Cosmides. "Groups in mind: The coalitional roots of war and morality." *Human Morality and Sociality: Evolutionary and Comparative Perspectives* (2010): 91-234.

Turner, John C. "Social comparison and social identity: Some prospects for intergroup behaviour." *European Journal of Social Psychology* 5, no. 1 (1975): 1-34.

Twitter. "About verified accounts." Help.twitter.com. https://help.twitter.com/en/managing-your-account/about-twitter-verified-accounts

Valentino, Nicholas A., Vincent L. Hutchings, Antoine J. Banks, and Anne K. Davis. "Is a worried citizen a good citizen? Emotions, political information seeking, and learning via the internet." *Political Psychology* 29, no. 2 (2008): 247-273.

Vance, J.D. ‘Life Outside the Liberal Bubble." New York Times, November 8, 2016. https://www.nytimes.com/interactive/projects/cp/opinion/election-night-2016/life-outside-the-liberal-bubble

Ward, Andrew, L. Ross, E. Reed, E. Turiel, and T. Brown. "Naive realism in everyday life: Implications for social conflict and misunderstanding." *Values and Knowledge* (1997): 103-135.

Watts, Duncan J., and Steven H. Strogatz. "Collective dynamics of ‘small-world’networks." *Nature* 393, no. 6684 (1998): 440-442.

Wayne, Teddy. "Clicking Their Way to Outrage." New York Times, July 3, 2014. https://www.nytimes.com/2014/07/06/fashion/social-media-some-susceptible-to-internet-outrage.html

Williams, Kipling D. "Ostracism." *Annual Review of Psychology* 58, no. 1 (2007): 425-452.

White, David Manning. "The “gate keeper”: A case study in the selection of news." *Journalism Bulletin* 27, no. 4 (1950): 383-390.

Wojcieszak, Magdalena E., and Diana C. Mutz. "Online groups and political discourse: Do online discussion spaces facilitate exposure to political disagreement?." *Journal of Communication* 59, no. 1 (2009): 40-56.

Wondra, Joshua D., and Phoebe C. Ellsworth. "An appraisal theory of empathy and other vicarious emotional experiences." *Psychological Review* 122, no. 3 (2015): 411-428.

Wall Street Journal. "Blue Feed, Red Feed." http://graphics.wsj.com/blue-feed-red-feed/

Zaller, John R. *The Nature and Origins of Mass Opinion*. Cambridge University Press, 1992.

Zhao, Shanyang, Sherri Grasmuck, and Jason Martin. "Identity construction on Facebook: Digital empowerment in anchored relationships." *Computers in Human Behavior* 24, no. 5 (2008): 1816-1836.

Zuckerberg, Mark. Facebook post. January 19, 2018, https://www.facebook.com/zuck/posts/10104445245963251

Zuckerman, Ethan. "Four problems for news and democracy." Medium. April 1, 2018, https://medium.com/trust-media-and-democracy/we-know-the-news-is-in-crisis-5d1c4fbf7691