Social Media, Echo Chambers, and Political Polarization

Pablo Barbera
University of Southern California
pbarbera@usc.edu

The unexpected rise of populist parties and candidates across developed democracies and the recent uptick in political violence in countries such as Myanmar, Sri Lanka or India have given urgency to the debate about the role that digital technologies and social media may be playing in exacerbating polarization and inciting extremist violence.

A popular argument that is commonly put forth as an explanation linking digital technologies to political polarization is related to their ability to foster the emergence of echo chambers where extremist ideas are amplified. Sunstein (2018), a leading proponent of this view, argues that the main characteristic of social networking sites is that they allow politically like-minded individuals to find one another. In this environment, citizens are only exposed to information that reinforces their political views and remain isolated from other individuals with opposing views, in part due to the filtering effects of ranking algorithms that generate filter bubbles (Pariser, 2011) and create incentives for publishers to share clickbait and hyperpartisan content (Benkler et al, 2018). The outcome of this process is a society that is increasingly segregated along partisan lines, and where compromise becomes unlikely due to rising mistrust on public officials, media outlets, and ordinary citizens on the other side of the ideological spectrum.

This view is now conventional wisdom not only among academics, but also in popular accounts of recent political developments. Just to give an example, former President Barack Obama referred to this argument in an interview with David Letterman:

> "If you are getting all your information off algorithms being sent through your phone and it's just reinforcing whatever biases you have, which is the pattern that develops, at a certain point, you just live in a bubble, and that's part of why our politics is so polarized right now. I think it's a solvable problem, but I think it's one we have to spend a lot of time thinking about."

> **Barack Obama**, *January 2018*, as quoted in CNN

Despite this apparent consensus, empirical studies offer a much more nuanced view of how social media affects political polarization, oftentimes questioning basic premises of this argument. Even if most political exchanges on social media take place among people with similar ideas, cross-cutting interactions are more frequent that commonly believed (Barbera, Jost, Nagler, Tucker and Bonneau, 2015), exposure to diverse news is higher than through other types of media (Barnidge, 2017; Fletcher and Nielsen, 2018a; Silver et al, 2019), and ranking algorithms don't have a large impact on the ideological balance of news consumption on

Facebook or Google (Bakshy, Messing and Adamic, 2015; Haim, Graefe and Brosius, 2018). A potential explanation for this series of findings is that social networking sites increase exposure to information shared by weak ties (Granovetter, 1977), such as co-workers, relatives, and acquaintances, who are more likely to share novel information – including ideologically diverse news (Bakshy et al., 2012; Barbera, 2015).

Of course, the fact that social media increases exposure to diverse political ideas from weak ties does not necessarily mean it has no effect on political polarization. Past research shows that repeated exposure to cross-cutting information leads to political moderation (Mutz, 2006), which could explain why political polarization in the United States has actually increased the least among those citizens who are least likely to use social media (Boxell, Gentzkow and Shapiro, 2017). However, a growing body of work challenges this finding, arguing that it is precisely this increased exposure to cross-cutting views that may be having polarizing effects (Bail et al., 2018; Suhay, Bello-Pardo and Maurer, 2018). In a lucid recent book, Settle (2018) identifies the heightened awareness of political identities in social media as a key factor driving affective polarization (Iyengar, Sood and Lelkes, 2012).

Despite the increased scholarly attention to this topic, much remains unknown. Most studies have focused their attention in the U.S. context and the comparative empirical evidence on this question is scarce, despite the well-known link between polarization and political violence in countries with high levels of ethnic fractionalization (see e.g. Montalvo and Reynal-Querol, 2005), which increases the urgency of research on this area. The lack of a shared definition of polarization makes it difficult to assess the overall impact of social media – the varying results across studies are likely due to differences in how polarization is conceptualized. We also know little about potential asymmetries in the polarizing effect of social media interactions regarding individuals' levels of political interest or their ideological inclinations.

The aim of this chapter is to offer an exhaustive review of the literature exploring the link between social media and political polarization. I highlight the areas where a consensus based on empirical evidence has already emerged, but also identify the questions that remain open, and the type of data and analysis that would help us address them.

Digital technologies and political echo chambers

The debate on whether digital technologies help bring politically different people together or create isolated ideological communities is almost as old as the internet. Possibly the earliest formulation can be attributed to Van Alstyne, Brynjolfsson et al. (1996), who warned that information technologies could lead to cyberbalkanization. Technology can "shrink geographic distances and facilitate information exchange," (p.3) thus creating a global village, but that was only one possible outcome. They hypothesized that an alternative scenario would be one where societies are more fragmented and interactions are balkanized because "the Internet makes it easier to find like-minded individuals," which "can facilitate and strengthen fringe communities that have a common ideology but are dispersed geographically" (p.5). And "once like-minded individuals locate each other," they conclude, "their subsequent interactions can further polarize their views or even ignite calls-to-action" (p.5).

This argument was further refined and popularized in Robert Putnam's seminal study of social capital in America, Bowling Alone (Putnam, 2000). In a remarkably prescient passage, Putnam

shares his concern about how "Internet technology" may allow and encourage "white supremacists to narrow their circle to like-minded intimates." While "real-world interactions often force us to deal with diversity," he continues, "the virtual world may be more homogeneous." This process could be accelerated by the use of "new 'filtering' technologies that automate the screening of 'irrelevant' messages" (p.178).

At the same time, Putnam also leaves room for alternative scenarios by arguing that "tendencies toward community homogeneity long predate the internet" and speculating that "weak ties that bridge among distinct groups might create an interwoven community of communities" (p.179).

Putnam's original formulation of this argument already outlined its three main components: (1) digital technologies facilitate the emergence of communities of like-minded individuals, where (2) they are increasingly isolated from any challenging information; a process that is exacerbated by (3) filtering algorithms. All three components are present in most subsequent studies in this question, often discussed within the context of other political communication theories. The following sections offer a more detailed discussion of each of these components, with an overview of the empirical evidence supporting or refuting them.

Communities of like-minded individuals

Access to the internet dramatically lowers the costs of exchanging messages and finding information regardless of geographic distance. Not being bound by physical proximity, citizens gain the ability to connect and organize based on shared interests, as niche or rare as they may be. These conversations may be taking place in the open, through open forums or blogs, but also in more private settings, such as closed groups on Facebook or private communities.

As Sunstein (2001) argues in Republic.com and its follow-up book, #Republic (Sunstein, 2018), online spaces create opportunities for enclave deliberation, which is the form of deliberation that takes place when conversations only occur among like-minded people. Enclave deliberation is not inherently negative. In fact, it can promote the development of positions that would otherwise be silenced or offer a safe space for people who suffer from discrimination. A recent prominent example is "Pantsuit Nation," a private Facebook group in favor of Hillary Clinton during the 2016 US presidential election campaign, where Clinton supporters – predominantly female – could express their opinions without fear of being harassed during a particularly polarizing election. There are many other cases of communities whose members share the same race or ethnic background, sexual orientation, interest in a particular issue, etc. that may not lead to undesirable outcomes from a normative perspective.

However, Sunstein (2001) argues that in practice the modal outcome of enclave deliberation is group polarization, which represents a "breeding ground for extremism" (p.71) and "may even put social stability at risk" (p.77). His concern here is that "through the mechanisms of social influence and persuasive arguments, members will move to positions that lack merit." This could happen either because the homogeneity of the group restricts the size of the argument pool or because individuals will be more likely to voice a popular opinion within the group in order to obtain the approval of as many members as possible. The existence of these two mechanisms has been demonstrated in a number of lab experiments that show that individuals who participate in homogeneous discussion groups tend to adopt more extreme positions after deliberating with their like-minded peers (see e.g. Myers and Lamm, 1976). In contrast, diverse groups tend to

perform better in solving problems even if the average ability of their members is lower (Hong and Page, 2004).

Sunstein argues that this mechanism is behind the rise in the number of hate groups and extremist organizations with a strong online presence, such as white supremacists and holocaust deniers. Other scholars have used this argument as an explanation for the emergence of groups with less nefarious intentions but whose success can be equally polarizing. For example, in their book "The Outrage Industry", Berry and Sobieraj (2013, p.165) document how digital tools were critical to the emergence and organization of the Tea Party movement. They lowered the barriers to entry for those who wanted to get involved and offered a space where political messages could be refined and agreed on.

The literature on social networks has also addressed this question in the context of studies of homophily; that is, individuals' propensity to cluster based on common traits, such as political ideology (McPherson, Smith-Lovin and Cook, 2001). Adamic and Glance (2005) analyzed the network of hyperlinks connecting the top 1,000 political blogs active during the 2004 US Presidential election, finding that "liberals and conservatives [link] primarily within their separate communities" (p.43). This applies not only to authors but also to blog readers: Lawrence, Sides and Farrell (2010) found that blog readers gravitate towards blogs that align with their political orientation and are more likely to be polarized than non-blog-readers.

While most of the early literature focused on internet in general, recent work has found similar patterns on social networking sites. Conover et al. (2012) showed that partisan Twitter users are significantly more likely to spread messages that are congenial with their ideological positions. According to Barbera, Jost, Nagler, Tucker and Bonneau (2015), close to 75% of retweets on political topics take place between users of similar ideological views. Del Vicario et al. (2016) found that information related to scientific news and conspiracy theories also tends to spread in homogeneous and polarized communities on Facebook. Moving beyond strictly political opinions, Aiello et al. (2012) showed that users with similar interests are more likely to be friends in online communities.

To put this set of results into context, it is important to note that these patterns may not necessarily generalize to all social media users. In fact, it is likely that the widespread perception of polarization on social media is due to a minority of highly active and visible partisan individuals. Barbera and Rivero (2015) offer evidence pointing in this direction: a minority of highly active individuals were responsible for the overwhelming majority of hyper-partisan content being spread on Twitter ahead of the 2012 U.S. presidential election. Similarly, a study of news sharing on Twitter (Shore, Baek and Dellarocas, 2016) found that only a small network core exhibited evidence of group polarization. Furthermore, as I discuss in the following section, the fact that users tend to self-select themselves into homogeneous communities does not necessarily imply that they are never exposed to dissonant political messages.

Avoiding opinion challenges?

The second component in the standard argument connecting digital technologies and political polarization is related to people's ability to filter out all information that may challenge their pre-existing views. As individuals increasingly spend their time in communities of like-minded individuals, they not only become more exposed to pro-attitudinal messages – their exposure to

counter-attitudinal information decreases as well. This is what leads to the emergence of echo chambers, where citizens do not see or hear a wide range of topics or ideas, which limits their capacity to reach common ground on political issues.

However, past work on selective exposure has found that reinforcement seeking and challenge avoidance are not intrinsically linked. Using survey data, Garrett (2009b) demonstrated that citizens take advantage of the broader availability of online news to increase exposure to political stories consistent with their views, but do not systematically avoid opinion challenges. In a follow-up study using behavior tracking in a realistic setting, Garrett (2009a) found no evidence that individuals abandon news stories that contain information with which they disagree once they have already started reading those stories.

Studies of individual-level online news consumption also offer scarce evidence that citizens actively avoid counter-attitudinal information. Gentzkow and Shapiro (2011) showed that ideological segregation in the news sites that people regularly visit is low in absolute terms, and not becoming higher over time. Similarly, Guess (2016) found high levels of overlap in news consumption: a majority of people mostly rely on centrist websites, and those who visit the most partisan websites are a minority that tends to be active news consumers and thus also visit many other sites.

This distinction between reinforcement seeking and challenge avoidance becomes relevant as we focus on social networking sites, where most of the political news that citizens consume are posted by their friends and family. The process whereby social boundaries weaken in online environments is what Brundidge (2010) refers to as the inadvertency thesis: somewhat counterintuitively, people are inadvertently exposed to additional political differences online. Furthermore, such exposure is not inconsequential. As Messing and Westwood (2014) demonstrate using a lab experiment that recreates a Facebook news feed, individuals are likely to click through and read news stories they potentially disagree with whenever those stories are recommended by their friends and acquaintances. In other words, friends' news recommendations are powerful social cues that can reduce partisan selective exposure to levels indistinguishable from chance.

Of course, this process of social consumption of information may still lead to echo-chamber environments if citizens' networks are politically homogeneous and thus all the content shared by social ties is pro-attitudinal. There is indeed some evidence that social ties in online networks replicate offline ties in their composition and nature (Jones et al., 2013; Bisbee and Larson, 2017). And yet, study after study demonstrates that despite citizens' propensity (and ability) to select their news sources from a broader range of sources thanks to digital technology, they are nonetheless exposed to a diverse set of views through websites and even more so on social media.

This finding is robust to the use of different sources of data and appears to hold across countries. Flaxman, Goel and Rao (2016) use behavioral data from web-browsing histories of 50,000 online adults that consume online news and offer probably the best evidence regarding the diverging patterns described above. On one hand, news consumption through social media and search engines contributes to the increased ideological segregation of audiences. However, at the individual level, these channels lead to more exposure to opposing perspectives than direct browsing. Similar to what Guess (2016) reports, this apparent paradox between macro- and

micro-level patterns can be explained by a minority of partisan individuals who are responsible for the majority of partisan news consumption.

Bakshy, Messing and Adamic (2015) report similar results on Facebook. Their analysis examines the ideological content of the Facebook news feeds of 10.1 million U.S. Facebook users. Unsurprisingly, they find that a majority of friendship links take place between people within the same ideological groups, and that users are indeed more likely to engage with congenial content. However, many friendships (20% for conservatives and 18% for liberals) do cut across ideological groups. And, most importantly, exposure to ideologically diverse news and opinion is also high: on average, around 30% of political news stories that users see in their news feed are cross-cutting. This proportion is remarkably similar to what Barbera (2015) reports using Twitter data: 33% of the tweets to which a sample of politically interested users were potentially exposed are cross-cutting.

Going beyond news consumption in an online setting, Barnidge (2017) offers a useful comparison of how U.S. adults report being exposed to political disagreement in different settings. His study relies on survey data which, at the expense of potential reporting biases, has the advantage of allowing to compare offline interactions and offline news consumption (newspapers, TV, etc.) with social media interactions. The results confirm the pattern described above: respondents admit being more exposed to political disagreement on social media than in other communication settings.

Although comparative studies of news consumption on social media are still scarce, the evidence presented by Fletcher and Nielsen (2017, 2018b) as part of the yearly Reuters Institute Digital News Report suggests that the U.S. case is not an exception. Using representative surveys conducted in 37 countries, they find that online audiences do not appear to be more politically fragmented than offline audiences, that incidental exposure to news content on social media is a generalized phenomenon, and that in fact people who use social networking sites are exposed to diverse news at a greater rate than those who do not. Similarly, a recent report by the Pew Research Center conducted in 11 nations across four global regions (Silver et al, 2019) found that social media users were more likely to interact regularly with a more diverse network, including people from different political parties, than those who were not active on social media.

How can we reconcile the fact that citizens now have a much greater ability to filter out any opinion challenges with the fact they actually do not appear to do so? One explanation, which I already advanced earlier, is that perhaps most people do not actually make an effort to avoid counter-attitudinal information (Garrett, 2009b), and in that sense the pattern we observe in an online setting replicates what we would observe offline. An alternative argument, however, is specific to social networking platforms – social media sites increase exposure to diverse views because they connect us to weak ties.

A majority of people with whom we interact in our daily life can be considered "weak ties," either because the frequency with which we communicate is low or because we don't perceive them to be close to us. In contrast, "strong ties" are that minority of people that we see frequently and whom we tend to trust more. Weak ties are generally co-workers, school friends, distant relative, and other acquaintances; whereas strong ties are our partners, close friends, and direct family. This distinction becomes relevant in the context of information diffusion because, as

Granovetter (1977) discovered, individuals are exposed to novel information (in his study, information about job opportunities) through weak ties.

When we think about social media sites, undoubtedly the main way in which they impact our daily lives is by making it easy to stay in touch with people we wouldn't see in person regularly. In other words, they entail greater exposure and contact with weak ties than in offline interactions (Gil de Zuniga and Valenzuela, 2011). As Gladwell (2010) argued in his popular New Yorker essay, "the platforms of social media are built around weak ties." His main point there is that such weak ties do not "lead to high-risk activism" – they do not generate the motivation that people need to protest. However, weak ties are fundamental if we want to understand the spread of information, because they help connect peripheral parts of online networks (Barbera, Wang, Bonneau, Jost, Nagler, Tucker and Gonzalez-Bailon, 2015; De Meo et al., 2014). That applies not only to protest movements but also to news consumption in general.

The other way in which tie strength is important to this argument is related to the fact that homophily tends to be lower among weak ties (McPherson, Smith-Lovin and Cook, 2001). In other words, because humans exhibit a propensity to preferentially establish links to other people that are similar to them, we should expect our weak ties to be more different from us than strong ties; and that applies also to ideological similarity. It is this heterogeneity that explains why weak ties are responsible for the propagation of novel information on social media (Bakshy et al., 2012).

To sum up, the most likely explanation for why citizens are not able to avoid opinion challenges on social media is the fact that (1) they don't have the ability to choose which content they see, since it's incidental and is presented in the context of powerful social cues, (2) most of that content is shared by weak ties, which tend to be more ideologically diverse than strong ties.

The filter bubble

The previous two sections dealt with two mechanisms at the core of how individuals choose to consume (or not) political information. However, in the current online environment, these are not the only two ingredients that determine citizens' media diet. As the number and heterogeneity of choices increase, citizens become unprepared to deal with the information overload that such wide availability entails. In this context, search engines and social networking sites have turned to rely on real-time, automated, and personalized algorithms to help users navigate their web browsing sessions.

In some cases, the set of heuristic and rules that determine how information is displayed on a website is well-understood. For example, search results on Google were ranked, at least in the early iterations of this website, using the PageRank algorithm developed by its founders (Page et al., 1999). However, when it comes to social media sites, the algorithms that determine how information is ranked on users' News Feed (on Facebook) or Timeline (on Twitter) are generally considered a black box. Other than some general principles about how these companies take a series of signals and try to make predictions about the relevancy score for each combination of user and piece of content (Mosseri, 2018), social networking sites have released little information about how this process really works. The most likely reason for this opaqueness is that releasing more information could make it easier for publishers that try to game the system.

However, this lack of transparency has also raised concerns about the extent to which these algorithms could actually be contributing to exacerbating inequalities and ideological segregation.

This concern is reminiscent of probably the earliest discussion of algorithmic curation of news, the "Daily Me" concept, popularized in the mid-1990s by MIT Media Lab founder Nicholas Negroponte (although the idea was conceived in the 1970s). In contrast to information consumed through newspapers and TV newscasts, where journalists act as the gatekeepers that determine which stories are newsworthy, the Daily Me would be tailored to each individual's preferences and interests. This "architecture machine" would "track what you skipped and reread and where you paused, then use those cues to evolve into a composite Daily Me that would carry only the news you cared about most" (Hapgood, 1995).

Sunstein (2018) uses this concept to justify his worry about the "age of the algorithm," in which citizens are not in control of the news they consume anymore. Even if people would not voluntarily choose to gravitate towards echo chambers, they may have no option, as social media sites become the arbiters of what people see and what they don't.

The most influential book within this line of thought is Eli Pariser's "The Filter Bubble" (Pariser, 2011). His concern applies to both search engines – which would yield completely different results based on the website's prediction of what the user's intention is – and social networking sites – which would show users only content that they are likely to prefer. As these online services learn more about the users, their accuracy in predicting what they would prefer to see becomes higher, eventually leading to bubbles where citizens would never be exposed to any type of information that they may cause them unpleasantness.

Pariser's worry is two-fold. First, he expresses concerns about algorithms increasing inequalities in civic engagement between politically interested people and those who would only prefer to consume entertainment news. And second, for those with clear partisan preferences, he claims that algorithms reduce the extent to which they "hear from the other side", even below the level of what they would voluntarily choose to see. The outcome of this dual process, which echoes similar concerns expressed by Prior (2007) and others about cable news, is a society where the type of shared experiences that are necessary in a well-functioning democracy are simply not non-existent.

Although this review of the existing work focuses on Pariser's influential manuscript, it is important to note that the concept of algorithmic bias has received much broader attention in political communication and computer science. Concerns about black-box deep learning methods, which sacrifice interpretability in order to improve their accuracy, have frequently been accused of having implicit biases. The fact that algorithms are automated does not imply that they don't replicate common human flaws, since the training dataset that they use to replicate human behavior in many cases also incorporate such biases (see e.g. Caliskan, Bryson and Narayanan, 2017).

In the specific case of ideological segregation, what do we know about the impact of ranking algorithms? The answer is... not much. There is a general lack of the type of systematic evidence that would allow us to answer key questions in the field. An explanation for this scarcity of evidence is researchers' inability to access or manipulate search or social ranking algorithms. It is thus not surprising that the most extensive study on this topic was conducted by Facebook

researchers Bakshy, Messing and Adamic (2015). Their results show that Facebook's algorithm induces a significant but small reduction in users' exposure to cross-cutting content. In contrast, "individual choices more than algorithms limit exposure to attitude-challenging content" (p.1131) on Facebook.

This study is not without limitations, however. As they acknowledge, their sample is only limited to active users who self-report their political affiliation on their profile. While this may lead to underestimating the extent to which users are exposed to cross-cutting views, when it comes to the relative importance of individual choices vs the algorithm, the opposite may be true. In other words, people who self-report their political views are more likely to be politically interested, and thus more likely to be more selective with the political content they consume. Individuals with lower levels of political interest may rely to a greater extent on Facebook's algorithm to decide which political news they are exposed to.

What about search engines, the other problematic source of news identified by Pariser (2011)? Again, the empirical evidence that goes beyond anecdotes is scarce. A recent study (Haim, Graefe and Brosius, 2018) casts doubts on the concerns about algorithmic filter bubbles. Using a set of different profiles in a realistic search scenario on Google, the authors found null effects on content diversity when experimenting with a set of politically heterogeneous profiles. These findings are consistent with results in Flaxman, Goel and Rao (2016) or Feuz, Fuller and Stalder (2011), which also found small effects of Google's search personalization algorithms on the information that users actually see. However, without access to Google's internal data and with search algorithms being a moving target given their constant evolution, it is hard to reach firm conclusions.

Social media and political polarization

The review of the literature on social media and "echo chambers" has shown that, rather counter-intuitively, there is convincing empirical evidence demonstrating that social networking sites increase the range of political views to which individuals are exposed. However, that tells us little regarding its subsequent impact on individual attitudes and aggregate-level polarization. As I describe next, there are good reasons to expect an effect in either direction.

Political theorists have long considered shared exposure to dissimilar views to be a necessary condition for the type of healthy political deliberation that takes place in thriving democratic societies (Habermas and Habermas, 1991; Mill, 1859). But beyond this normative argument, diverse deliberation is also important because it can have a profound impact on citizens' beliefs – and their strength. Cross-cutting interactions increase awareness of rationales for individuals' own viewpoints and oppositional viewpoints (Mutz, 2006). Exposure to diverse opinions can also be framed within the "contact hypothesis" literature. Based on Allport's (1954) classic intergroup hypothesis, this argument would imply that exposure to members of different groups (in this case, people who support a different party) should reduce prejudice and promote political tolerance.

If social media use indeed increases citizens' awareness of diverse viewpoints and fosters intergroup contact, it seems reasonable to expect that it may also weaken the strength of people's political beliefs and thus reduce political polarization. Although a review of how these two mechanisms operate in the offline context is beyond the scope of this chapter, there is indeed

broad evidence that cross-cutting interactions foster political moderation (Carpini, Cook and Jacobs, 2004), e.g. by increasing support for civil liberties of disliked groups (Mutz, 2002), and that intergroup contact reduces prejudice (Pettigrew and Tropp, 2006; Paluck, Green and Green, 2018).

However, the empirical studies that examine this question in the social media context are few and offer mixed results. Some studies find a negative link between social media usage and political polarization, as the literature discussed above would lead us to expect. Probably the most prominent example is the work by Boxell, Gentzkow and Shapiro (2017), who find that even if the level of political polarization has increased for all age cohorts in the U.S., the change has been smallest in magnitude among the youngest set of people – those most likely to be active on social media. This result suggests that digital technologies are likely to play a limited role in explaining why polarization is on the rise. One important limitation of this study, however, is that its longitudinal analysis of age cohorts could suffer from ecological inference problems if, for example, their composition in terms of unobserved demographic covariates has changed over time.

Despite any concerns in this particular article, at least three other studies using individual-level data yield results pointing in the same direction. Heatherly, Lu and Lee (2017) find that individuals who engage in cross-cutting discussion on social media report lower levels of political polarization. Barbera (2015) finds that individuals in ideologically diverse Twitter networks tend to moderate their ideological positions over time. Using survey data from 28 European countries, Ngyuen and Vu (2019) find that citizens who consume political news through social media sites are not more polarized than those who rely on other sources.

Other studies, on the other hand, do suggest that exposure to political information through social media could have polarizing effects. Bail et al. (2018) show that cross-cutting exposure to political messages from elites can increase polarization. Using an innovative research design that maximizes internal and external validity, the authors recruited a sample of respondents and then asked them to follow bots that were sharing political messages that were counter-attitudinal with respect to their own views. A longitudinal comparison of the respondents' political opinions showed that Republicans became significantly more conservative. Democrats also became somewhat more liberal, although the change was not statistically significant. These results offer evidence of a backlash effect that could be due to motivated reasoning (Lodge and Taber, 2013), perhaps operating in a similar way as the boomerang effects that have been found in some fact-checking interventions to rumors and misinformation (Nyhan and Reifler, 2010). Note, however, that the existence of these backfire effects is highly debated, as summarized in by Wittenberg and Berinsky (this volume).

Similarly, two experimental studies conducted by Suhay, Bello-Pardo and Maurer (2018) also find that exposure to political disagreement in online settings increases political polarization. These disagreements are presented in an uncivil context which, in their view, is representative of the type of cross-cutting interactions that take place online. It is the combination of both criticism of partisan identities, and not necessarily online opinions on specific issues, that drives polarization.

One last article that examined the overall polarizing effects of social media is Allcott et al (2019). The authors used an incentive design to study how leaving Facebook affects a variety of

outcomes, including political polarization. The results reveal that deactivation did not significantly move affective polarization or polarization in factual beliefs about current events; but it reduced polarization of views on policy issues, using a scale similar to Bail et al (2018). However, subjects' knowledge about current events also decreased, which suggests that the depolarizing effects of leaving social media may be explained by a lower level of exposure to political information overall. An additional limitation of this study is that it focused on the effects of stopping using social media, which may not necessarily be symmetrical to how the adoption of these platforms affects individual-level attitudes and behavior.

In conclusion, we see that the existing literature offers results that appear to be at odds. How can we reconcile these diverging findings? To answer this question, it is important to consider how we define political polarization – do we mean just divergence in political views or issue positions (ideological polarization) or dislike for the partisan outgroup (affective polarization)? In addition, it is also plausible that the differences across empirical studies could be due to social media having a heterogeneous effects across different groups of people, particularly regarding their political orientation and the strength of their partisan identities.

Ideological polarization and affective polarization

We generally think of information environments where individuals are exposed to multiple viewpoints as spaces that should lead to social consensus (DeGroot, 1974). However, this argument assumes that individuals do not experience cognitive biases in how they process the messages they receive, and that they are Bayesian learners, updating in the direction of the information they receive. Research on partisan motivated reasoning (Lodge and Taber, 2013) shows that, when it comes to political information, that may not be the case. For example, when presented with negative information about a candidate that citizens evaluate positively, support for that candidate may actually increase due to affective biases (Redlawsk, 2002).

Biased information processing is a relevant factor driving the perceived increase in polarization among the American public. Even if Democrats and Republicans may agree on some policy positions, they tend to increasingly dislike and distrust one another, and to perceive that the social distance between them has expanded (Mason, 2018). This alternative conceptualization of polarization is often called "affective polarization" (Iyengar, Sood and Lelkes, 2012) or "psychological polarization" (Settle, 2018) as opposed to "ideological polarization," which would be limited to divergence of political views. This distinction may help us reconcile the mixed empirical results: it is possible that increased exposure to cross-cutting views on social media is moderating citizens' political views while at the same time exacerbating their perceived social distance with respect to the partisan outgroup.

Settle (2018) articulates this argument in her recent book, "Frenemies. How Social Media Polarizes America." She convincingly argues that core features of social media platforms, such as the fusion of social and political content, the ubiquity of social feedback, the ability to easily infer other users' political identity or the incentives to produce "clickbait-y" and inflammatory content to catch people's attention, have a direct impact on the aggregate level of psychological polarization. The mechanism behind these effects is that these features facilitate the type of psychological processes behind affective polarization: the reinforcement of social and political identities, in combination with citizens' biased information processing. As Settle (2018) shows through a series of survey studies and lab experiments, social media usage (and the type of

partisan content that is shared on Facebook) increases the perceived differences between individuals' own position and where they perceive the outgroup to be, makes political and social identities more correlated, and contributes to the stereotyping (and negative evaluations) of the outgroup.

Drawing on the distinction between ideological and affective polarization helps us reconcile the diverging empirical findings but also allows us to derive useful normative implications. Although ideological polarization is generally considered to be undesirable, its absence could also counter-productive. It can be a sign that political competition is not clearly structured and that parties do not offer distinct positions, which could depress civic engagement. That was actually the concern behind the 1950 report "Toward a More Responsible Two-Party System," developed by the American Political Science Association (APSA) Committee on Political Parties, which has received increased attention recently (Wickham-Jones, 2018). In contrast, the normative debate about the pernicious effects of affective polarization is settled: when the perceived distance between partisan groups is large, debates are won based on identities and not on substantive arguments and empirical evidence, which in turn can have negative economic consequences such as increased inequality and unemployment (Bonica et al., 2013; Enns and McAvoy, 2012; Mason, 2018).

Asymmetric polarization

There are well-known differences in how individuals select, consume, and process political information depending on their political orientation, interest in politics, and strength of partisanship (see e.g. Zaller et al., 1992; Graber, 1988; Prior, 2007). For that reason, expecting to find that social media usage has a homogeneous (de)polarizing effect for all citizens seems too simplistic; in the same way that not all individuals are equally susceptible to misinformation (see Wittenberg and Berinsky, this volume). However, citizens' online experiences are shaped in predictable ways by their own political and sociodemographic characteristics. Understanding how different research designs – in particular sampling decisions – capture these factors may also help us explain the diverging results described in the previous sections. Here I offer an overview of research that suggests that two of them – political interest and ideological orientation – may be playing such a role.

In his classic study on the origins of public opinion, Zaller et al. (1992) argues that politically aware individuals are more receptive to pro-attitudinal messages. Similarly, Taber and Lodge (2006) find that those with highest levels of political sophistication are more likely to uncritically accept supporting arguments and reject counter-attitudinal arguments, leading to attitude polarization. And just like beliefs in misinformation are more difficult to correct among the most informed citizens (see Wittenberg and Berinsky, this volume), deactivating polarization among this same subset of the population may require more than just increasing exposure to cross-cutting political views. This pattern could be exacerbated simply by citizens' self-selection into the consumption of news vs entertainment, based on their pre-existing level of political interest (Prior, 2007; Stroud, 2010).

These are thus good reasons to expect that social media interactions may lead to polarization among the minority of partisan individuals who are most active in discussions about politics on social media (Benkler et al, 2018; Barbera and Rivero, 2015; Shore, Baek and Dellarocas, 2016). This may help explain the backlash reaction to cross-cutting interactions identified by Bail

et al. (2018), whose sample contain only self-identified partisans. It is also consistent with the high degrees of ideological segregation found by Conover et al. (2012) or Halberstam and Knight (2016), whose analyses focuses on politically engaged Twitter users.

Political ideology – or personality traits that have been found to be correlated with it – is a final factor that can help us understand polarization patterns on social media. Examining consumption of traditional media, some past research has found that conservatives are more likely than liberals to engage in selective exposure, biased information processing, and ideological conformity (Lau and Redlawsk, 2006; Nam, Jost and Van Bavel, 2013; Nyhan and Reifler, 2010; Garrett, 2009b; Guess et al, 2019), although other work has found symmetric patterns regarding these behaviors (Iyengar and Hahn, 2009; Munro et al., 2002; Nisbet, Cooper and Garrett, 2015).

When it comes to social media more specifically, Barbera, Jost, Nagler, Tucker and Bonneau (2015) showed that these asymmetries extend to ideological segregation in information diffusion on Twitter. Their analysis found that liberals were more likely to engage in cross-ideological interactions than conservatives. This finding is consistent with the results of the field experiment conducted by Bail et al. (2018), which found that the increase in polarization in reaction to counter-attitudinal information was larger for conservatives than for liberals. It is possible, however, that some of these results are confounded by personality traits that some researchers consider to be correlated with ideology, such as openness to experiences and conscientiousness (Carney et al., 2008). For example, Kim, Hsu and de Zuniga (2013) show that social media usage has a smaller effect on the overall heterogeneity of discussion networks among conservatives than liberals. This finding needs to be understood within the broader media ecosystem in the United States which, as Benkler et al (2018) show, exhibits a similar asymmetry in terms of audience and production, with political clickbait and hyperpartisan content being more prevalent on the right than on the left.

The path forward: what we don't know

The research summarized in this chapter has greatly advanced our understanding of how the success of social networking platforms is transforming the structure and heterogeneity of political conversations, and its subsequent impact on political polarization. As we have seen, in many cases the empirical evidence has challenged the conventional wisdom, while in other cases the theory has moved ahead of the evidence and can help us reconcile seemingly contradictory findings. Although the field is reaching a state of maturity, many questions remain still open. This concluding section offers a detailed list of what we don't know (yet).

Do the results described here generalize to other contexts beyond the U.S.? Mirroring the state of the literature on political polarization more generally, most of what we know about cyberbalkanization is based on data from the U.S. only. With some notable exceptions (Bechmann and Nielbo, 2018; Bright, 2018; Vaccari et al., 2016), there is a lack of theory and empirical evidence about how contextual factors mediate the relationship between social media usage and political polarization from a comparative perspective. And while generally we think about polarization in terms of the gridlock and suboptimal economic outcomes it can generate, in some instances it can lead to much worse outcomes, including political violence. For example, a working paper has linked hate speech spread on social media to anti-refugee attacks in Germany (Muller and Schwarz, 2017). Journalistic accounts of religious violence in Sri Lanka and

Myanmar have shown how these events appear to be fueled by rumors spread on Facebook and Whatsapp (Taub and Fisher, 2018). These two instances illustrate the urgent need for research on how the spread of extremist ideas on social media could be motivating offline violence.

Is there any longitudinal variation regarding the polarizing effects of social media interactions? Are things getting better or worse? In many ways, the study of how digital technologies affect political behavior is a moving target (Munger, 2019). Internet services are in constant evolution, both in terms of which social media platforms or websites become popular and regarding the features of those platforms. Even if Facebook has now existed for more than 10 years, it has significantly evolved over this period. How the platforms are used can also change: Twitter was initially just a website where users would post real-time updates about their everyday activities. In 2018, it has become the go-to place for breaking news. The population that is active on these platforms can also radically change: in 2014, there were 2 million Facebook users in Myanmar; today, that number has skyrocketed to more than 30 million (Roose, 2018). It is true that some of the findings in this chapter refer to mechanisms that are a core component of human behavior – for example, biased information processing or the structure of social networks. However, studying a constantly shifting environment may mean an inability to understand the scope conditions under which those mechanisms hold.

How can we disentangle individuals' self-selection into networks from the impact on their attitudes? Most the empirical research cited in this paper relies on cross-sectional evidence from observational studies. One challenge when deriving causally valid conclusions from this evidence is that individuals' media consumption and exposure to political messages on social media is highly endogenous. For example, when we find that individuals in heterogeneous networks are less polarized, an alternative interpretation is that individuals with moderate political positions are more likely to feel comfortable when they are exposed to cross-cutting political views (Barbera, 2015; Kim, Hsu and de Zuniga, 2013; Vaccari et al., 2016). Disentangling this reverse causality loop is key in order to make progress within this field. In that sense, Bail et al. (2018) and Allcott et al (2019) offer a blueprint to imitate: a large-scale field experiment where exposure to social media messages is randomly assigned, which allows the authors to make valid causal claims. Ethical considerations should also be part of this debate, however, particularly as scholars focus their efforts on studying how extremism fueled by social media interactions can lead to offline violence.

What are the unintended consequences of potential interventions to reduce polarization? Concerns about the negative societal consequences of political polarization may lead to regulatory changes or new platform features to deactivate it. For example, Settle (2018, ch.9) offers a helpful list of suggestions that include increasing information quality on social media profiles, incentivizing moderation by upranking reasonable disagreement, and eliminating highly visible quantification. As we consider the benefits and disadvantages of these options, we should bear in mind the trade-offs that fostering moderation involves. The work by Mutz (2002) is a perfect illustration of these challenges: exposure to civil, cross-cutting exchanges may lead to political tolerance, but it can also make politics more complex and less interesting, which depresses civic engagement and accentuates political inequality. Studying these complex multi-causal relationships is essential in order to address one of the most important questions of our time – how digital technologies affect democratic politics.

References

Adamic, Lada A and Natalie Glance. 2005. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*. ACM pp. 36–43.

Aiello, Luca Maria, Alain Barrat, Rossano Schifanella, Ciro Cattuto, Benjamin Markines and Filippo Menczer. 2012. "Friendship prediction and homophily in social media." *ACM Transactions on the Web (TWEB)* 6(2):9.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2019. *The welfare effects of social media*. NBER Working Paper No. 25514.

Allport, Gordon Willard. 1954. "The nature of prejudice.".

Bail, Christopher A, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout and Alexander Volfovsky. 2018. "Exposure to opposing views on social media can increase political polarization." *Proceedings of the National Academy of Sciences* p. 201804840.

Bakshy, Eytan, Itamar Rosenn, Cameron Marlow and Lada Adamic. 2012. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*. ACM pp. 519–528.

Bakshy, Eytan, Solomon Messing and Lada A Adamic. 2015. "Exposure to ideologically diverse news and opinion on Facebook." *Science* 348(6239):1130–1132.

Barbera, Pablo. 2015. "How social media reduces mass political polarization. Evidence from Germany, Spain, and the US." *Paper presented at the 2015 APSA conference* .

Barbera, Pablo and Gonzalo Rivero. 2015. "Understanding the political representativeness of Twitter users." *Social Science Computer Review* 33(6):712–729.

Barbera, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker and Richard Bonneau. 2015. "Tweeting from left to right: Is online political communication more than an echo chamber?" *Psychological science* 26(10):1531–1542.

Barbera, Pablo, Ning Wang, Richard Bonneau, John T Jost, Jonathan Nagler, Joshua Tucker and Sandra Gonzalez-Bailon. 2015. "The critical periphery in the growth of social protests." *PloS one* 10(11):e0143611.

Barnidge, Matthew. 2017. "Exposure to political disagreement in social media versus face-to-face and anonymous online settings." *Political Communication* 34(2):302–321.

Bechmann, Anja and Kristoffer L Nielbo. 2018. "Are We Exposed to the Same News in the News Feed? An Empirical Analysis of Filter Bubbles as Information Similarity for Danish Facebook Users." *Digital Journalism* pp. 1–13.

Benkler, Yochai, Robert Faris, and Hal Roberts. 2018. *Network Propaganda. Manipulation, Disinformation, and Radicalization in American Politics*. Oxford University Press.

Berry, Jeffrey M and Sarah Sobieraj. 2013. *The outrage industry: Political opinion media and the new incivility.* Oxford University Press.

Bisbee, James and Jennifer M Larson. 2017. "Testing Social Science Network Theories with Online Network Data: An Evaluation of External Validity." *American Political Science Review* 111(3):502–521.

Bonica, Adam, Nolan McCarty, Keith T Poole and Howard Rosenthal. 2013. "Why hasn't democracy slowed rising inequality?" *Journal of Economic Perspectives* 27(3):103–24.

Boxell, Levi, Matthew Gentzkow and Jesse M Shapiro. 2017. "Greater Internet use is not associated with faster growth in political polarization among US demographic groups." *Proceedings of the National Academy of Sciences* p. 201706588.

Bright, Jonathan. 2018. "Explaining the emergence of political fragmentation on social media: the role of ideology and extremism." *Journal of Computer-Mediated Communication* 23(1):17–33.

Brundidge, Jennifer. 2010. "Encountering difference in the contemporary public sphere: The contribution of the Internet to the heterogeneity of political discussion networks." *Journal of Communication* 60(4):680–700.

Caliskan, Aylin, Joanna J Bryson and Arvind Narayanan. 2017. "Semantics derived automatically from language corpora contain human-like biases." *Science* 356(6334):183–186.

Carney, Dana R, John T Jost, Samuel D Gosling and Jeff Potter. 2008. "The secret lives of liberals and conservatives: Personality profiles, interaction styles, and the things they leave behind." *Political Psychology* 29(6):807–840.

Carpini, Michael X Delli, Fay Lomax Cook and Lawrence R Jacobs. 2004. "Public deliberation, discursive participation, and citizen engagement: A review of the empirical literature." *Annu. Rev. Polit. Sci.* 7:315–344.

Conover, Michael D, Bruno Goncalves, Alessandro Flammini and Filippo Menczer. 2012. "Partisan asymmetries in online political activity." *EPJ Data Science* 1(1):6.

De Meo, Pasquale, Emilio Ferrara, Giacomo Fiumara and Alessandro Provetti. 2014. "On Facebook, most ties are weak." *Communications of the ACM* 57(11):78–84.

DeGroot, Morris H. 1974. "Reaching a consensus." *Journal of the American Statistical Association* 69(345):118–121.

Del Vicario, Michela, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley and Walter Quattrociocchi. 2016. "The spreading of misinformation online." *Proceedings of the National Academy of Sciences* 113(3):554–559.

Enns, Peter K and Gregory E McAvoy. 2012. "The role of partisanship in aggregate opinion." *Political Behavior* 34(4):627–651.

Feuz, Martin, Matthew Fuller and Felix Stalder. 2011. "Personal Web searching in the age of semantic capitalism: Diagnosing the mechanisms of personalisation." *First Monday* 16(2).

Flaxman, Seth, Sharad Goel and Justin M Rao. 2016. "Filter bubbles, echo chambers, and online news consumption." *Public opinion quarterly* 80(S1):298–320.

Fletcher, Richard and Rasmus Kleis Nielsen. 2017. "Are news audiences increasingly fragmented? a cross-national comparative analysis of cross-platform news audience fragmentation and duplication." *Journal of Communication* 67(4):476–498.

Fletcher, Richard and Rasmus Kleis Nielsen. 2018a. "Are people incidentally exposed to news on social media? A comparative analysis." *New Media & Society* 20(7):2450–2468.

Fletcher, Richard and Rasmus Kleis Nielsen. 2018b. "Are people incidentally exposed to news on social media? A comparative analysis." *New media & society* 20(7):2450–2468.

Garrett, R Kelly. 2009a. "Echo chambers online?: Politically motivated selective exposure among Internet news users." *Journal of Computer-Mediated Communication* 14(2):265–285.

Garrett, R Kelly. 2009b. "Politically motivated reinforcement seeking: Reframing the selective exposure debate." *Journal of Communication* 59(4):676–699.

Gentzkow, Matthew and Jesse M Shapiro. 2011. "Ideological segregation online and offline." *The Quarterly Journal of Economics* 126(4):1799–1839.

Gil de Zuniga, Homero and Sebastian Valenzuela. 2011. "The mediating path to a stronger citizenship: Online and offline networks, weak ties, and civic engagement." *Communication Research* 38(3):397–421.

Gladwell, Malcolm. 2010. "Small change." *The New Yorker*.

Graber, Doris Appel. 1988. *Processing the news: How people tame the information tide.* University Press of Amer.

Granovetter, Mark S. 1977. The strength of weak ties. In *Social networks*. Elsevier pp. 347–367.

Guess, Andrew M. 2016. "Media choice and moderation: Evidence from online tracking data." *Unpublished manuscript, New York University.*

Guess, Andrew M., Jonathan Nagler, and Joshua Tucker. 2019. "Less than you think: Prevalence and predictors of fake news dissemination on Facebook". *Science Advances*, 5(1): eaau4586.

Habermas, Jurgen and Jurgen Habermas. 1991. *The structural transformation of the public sphere: An inquiry into a category of bourgeois society.* MIT press.

Haim, Mario, Andreas Graefe and Hans-Bernd Brosius. 2018. "Burst of the Filter Bubble?" *Digital Journalism* (3):330–343.

Halberstam, Yosh and Brian Knight. 2016. "Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter." *Journal of Public Economics* 143:73–88.

Hapgood, Fred. 1995. "The Media Lab at 10." *Wired Magazine* .

Heatherly, Kyle A, Yanqin Lu and Jae Kook Lee. 2017. "Filtering out the other side? Cross-cutting and like-minded discussions on social networking sites." *New media & society* 19(8):1271–1289.

Hong, Lu and Scott E Page. 2004. "Groups of diverse problem solvers can outperform groups of high-ability problem solvers." *Proceedings of the National Academy of Sciences* 101(46):16385–16389.

Iyengar, Shanto, Gaurav Sood and Yphtach Lelkes. 2012. "Affect, not ideologya social identity perspective on polarization." *Public opinion quarterly* 76(3):405–431.

Iyengar, Shanto and Kyu S Hahn. 2009. "Red media, blue media: Evidence of ideological selectivity in media use." *Journal of Communication* 59(1):19–39.

Jones, Jason J, Jaime E Settle, Robert M Bond, Christopher J Fariss, Cameron Marlow and James H Fowler. 2013. "Inferring tie strength from online directed behavior." *PloS one* 8(1):e52168.

Kim, Yonghwan, Shih-Hsien Hsu and Homero Gil de Zuniga. 2013. "Influence of social media use on discussion network heterogeneity and civic engagement: The moderating role of personality traits." *Journal of Communication* 63(3):498–516.

Lau, Richard R and David P Redlawsk. 2006. *How voters decide: Information processing in election campaigns.* Cambridge University Press.

Lawrence, Eric, John Sides and Henry Farrell. 2010. "Self-segregation or deliberation? Blog readership, participation, and polarization in American politics." *Perspectives on Politics* 8(1):141–157.

Lodge, Milton and Charles S Taber. 2013. *The rationalizing voter*. Cambridge University Press.

Mason, Lilliana. 2018. *Uncivil agreement: How politics became our identity*. University of Chicago Press.

McPherson, Miller, Lynn Smith-Lovin and James M Cook. 2001. "Birds of a feather: Homophily in social networks." *Annual review of sociology* 27(1):415–444.

Messing, Solomon and Sean J Westwood. 2014. "Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online." *Communication Research* 41(8):1042–1063.

Mill, John Stuart. 1859. *On Liberty, Utilitarianism, and other essays.*

Montalvo, Jose G and Marta Reynal-Querol. 2005. "Ethnic polarization, potential conflict, and civil wars." *American economic review* 95(3):796–816.

Mosseri, Adam. 2018. "News Feed Ranking in Three Minutes Flat.".
URL: https://newsroom.fb.com/news/2018/05/inside-feed-news-feed-ranking/

Muller, Karsten and Carlo Schwarz. 2017. "Fanning the Flames of Hate: Social Media and Hate Crime.".

Munger, Kevin. 2019. "Knowledge Decays: Temporal Validity and Social Science in a Changing World". *Working paper*, https://osf.io/4utsk/

Munro, Geoffrey D, Peter H Ditto, Lisa K Lockhart, Angela Fagerlin, Mitchell Gready and Elizabeth Peterson. 2002. "Biased assimilation of sociopolitical arguments: Evaluating the 1996 US presidential debate." *Basic and Applied Social Psychology* 24(1):15–26.

Mutz, Diana C. 2002. "Cross-cutting social networks: Testing democratic theory in practice." *American Political Science Review* 96(1):111–126.

Mutz, Diana C. 2006. *Hearing the other side: Deliberative versus participatory democracy.* Cambridge University Press.

Myers, David G and Helmut Lamm. 1976. "The group polarization phenomenon." *Psychological bulletin* 83(4):602.

Nam, H Hannah, John T Jost and Jay J Van Bavel. 2013. "Not for all the tea in China! Political ideology and the avoidance of dissonance-arousing situations." *PloS one* 8(4):e59837.

Nguyen, An, and Hong Tien Vu. 2019. "Testing popular news discourse on the "echo chamber" effect: Does political polarisation occur among those relying on social media as their primary politics news source?." *First Monday* 24 (5).

Nisbet, Erik C, Kathryn E Cooper and R Kelly Garrett. 2015. "The partisan brain: How dissonant science messages lead conservatives and liberals to (dis) trust science." *The ANNALS of the American Academy of Political and Social Science* 658(1):36–66.

Nyhan, Brendan and Jason Reifler. 2010. "When corrections fail: The persistence of political misperceptions." *Political Behavior* 32(2):303–330.

Page, Lawrence, Sergey Brin, Rajeev Motwani and Terry Winograd. 1999. The PageRank citation ranking: Bringing order to the web. *Technical report Stanford InfoLab.*

Paluck, Elizabeth Levy, Seth A Green and Donald P Green. 2018. "The contact hypothesis re-evaluated." *Behavioural Public Policy* pp. 1–30.

Pariser, Eli. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.

Pettigrew, Thomas F and Linda R Tropp. 2006. "A meta-analytic test of intergroup contact theory." *Journal of personality and social psychology* 90(5):751.

Prior, Markus. 2007. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press.

Putnam, Robert D. 2000. *Bowling alone: America's declining social capital*. Simon and Schuster.

Redlawsk, David P. 2002. "Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making." *The Journal of Politics* 64(4):1021–1044.

Roose, Kevin. 2018. "Forget Washington. Facebooks Problems Abroad Are Far More Disturbing.".
URL: https://www.nytimes.com/2017/10/29/business/facebook-misinformation-abroad.html

Settle, Jaime E. 2018. Frenemies: *How Social Media Polarizes America*. Cambridge University Press.

Silver, Laura, Christine Huang and Kyle Taylor. 2019, "In Emerging Economies, Smartphone and Social Media Users Have Broader Social Networks". *Pew Research Center Report*.

Shore, Jesse, Jiye Baek and Chrysanthos Dellarocas. 2016. "Network structure and patterns of information diversity on Twitter." *arXiv preprint arXiv:1607.06795* .

Stroud, Natalie Jomini. 2010. "Polarization and partisan selective exposure." *Journal of communication* 60(3):556–576.

Suhay, Elizabeth, Emily Bello-Pardo and Brianna Maurer. 2018. "The polarizing effects of online partisan criticism: evidence from two experiments." *The International Journal of Press/Politics* 23(1):95–115.

Sunstein, Cass R. 2001. *Republic.com*. Princeton University Press.

Sunstein, Cass R. 2018. *#Republic: Divided democracy in the age of social media.* Princeton University Press.

Taber, Charles S and Milton Lodge. 2006. "Motivated skepticism in the evaluation of political beliefs." *American Journal of Political Science* 50(3):755–769.

Taub, Amanda and Max Fisher. 2018. "Where Countries are Tinderboxes and Facebook is a match." *New York Times*, April 21.

Vaccari, Cristian, Augusto Valeriani, Pablo Barbera, John T Jost, Jonathan Nagler and Joshua A Tucker. 2016. "Of echo chambers and contrarian clubs: exposure to political disagreement among German and Italian users of Twitter." *Social Media+ Society* 2(3):2056305116664221.

Van Alstyne, Marshall, Erik Brynjolfsson et al. 1996. Electronic communities: Global village or cyberbalkans. In *Proceedings of the 17th International Conference on Information Systems*. New York: Wiley. p. 32.

Wickham-Jones, Mark. 2018. "This 1950 political science report keeps popping up in the news. Heres the story behind it.".
URL: https://www.washingtonpost.com/news/monkey-cage/wp/2018/07/24/this-1950-political-science-report-keeps-popping-up-in-the-news-heres-the-story-behind-it/

Zaller, John R et al. 1992. *The nature and origins of mass opinion*. Cambridge university press.