

#republic: Divided Democracy in the Age of Social Media (2017)

Cass Sunstein

THE DAILY ME

[1] In 1995, MIT technology specialist Nicholas Negroponte prophesied the emergence of "*the Daily Me*." With the *Daily Me*, he suggested, you would not rely on the local newspaper to curate what you saw, and you could bypass the television networks. Instead, you could design a communications package just for you, with each component fully chosen in advance. If you want to focus only on basketball, you could do exactly that. If your taste runs to William Shakespeare, your *Daily Me* could be all Shakespeare, all the time. If you like to read about romances— perhaps involving your favorite celebrities— your newspaper could focus on the latest love affairs, or who's breaking up with whom. Or suppose that you have a distinctive point of view. Maybe your views are left of center, and you want to read stories fitting with what you think about climate change, equality, immigration, and the rights of labor unions. Or maybe you lean to the right, and you want to see conservative perspectives on those issues, or maybe on just one or two, and on how to cut taxes and regulation, or reduce immigration.

[2] Perhaps what matters most to you are your religious convictions, and you want to read and see material with a religious slant (your own). Perhaps you want to speak to and hear from your friends, who mostly think as you do; you might hope that all of you will share the same material. What matters is that with the *Daily Me*, everyone could enjoy an architecture of control. Each of us would be fully in charge of what we see and hear. In countless domains, human beings show "**homophily**": a **strong tendency to connect and bond with people who are like them**. The tendency to homophily is dampened if people live within social architectures that expose them to diverse types of people— in terms of perspectives, interests, and convictions. But with an architecture of control, birds of a feather can easily flock together.

[3] In the 1990s, the idea of a *Daily Me* seemed more than a little absurd. But it's looking astoundingly good. If anything, Negroponte understated what was coming, what has now arrived, and what is on the horizon. Is that a promise or a threat? I think it's both— and that the threatening part is what needs to be emphasized, not least because so many people see it as pure promise. True, there's no *Daily Me*, at least not quite yet. But we're getting there. Most Americans now receive much of their news from social media, and all over the world, Facebook has become central to people's experience of the world. It used to be said that the "Revolution Will Not Be Televised"; maybe or maybe not, but you can be pretty sure that the revolution will be tweeted (#Revolution). In 2016, for example, the military attempted a coup in Turkey. It succeeded in seizing the nation's major television network. But it failed to take over social media, which the government and its supporters successfully used to call the public to the streets and, in short order, to stabilize the situation. Coup attempts often stand or fall on public perceptions of whether they are succeeding, and social media played a major role in combating the perception that the government was falling.

[4] When people use Facebook to see exactly what they want to see, their understanding of the world can be greatly affected. Your Facebook friends might provide a big chunk of the news on which you focus, and if they have a distinctive point of view, that's the point of view that you'll see most. I worked in the Obama administration, and so did a number of my Facebook friends, and what I see on my Facebook page often fits the interests and views of the kind of people who worked in the Obama administration. Is that an unalloyed good? Probably not. And I have conservative friends whose Facebook pages look radically different from mine, in ways that fit with their political convictions. We are living in different political universes— something like science fiction's parallel worlds. A lot of the supposed news is fake. Your Twitter feed might well reflect your preferred topics and convictions, and it might provide much of what you see about politics— taxes, immigration, civil rights, and war and peace. What comes in your feed is your choice, not anyone else's. You might well choose to include topics that interest you, and points of view that you find congenial. In fact that seems quite natural. Why would you want topics that bore you and perspectives that you despise?

ALGORITHMS AND HASHTAGS

[5] As it turns out, you do not need to create a *Daily Me*. Others are creating it for you right now (and you may have no idea that they're doing it). Facebook itself does some curating, and so does Google. We live in the age of the algorithm, and the algorithm knows a lot. With the rise of artificial intelligence, algorithms are bound to improve immeasurably. They will learn a great deal about you, and they will know what you want or will like, before you do, and better than you do. They will even know your emotions, again before and better than you do, and they will be able to mimic emotions on their own. Even now, an algorithm that learns a little bit about you can discover and tell you what "people like you" tend to like. It can create something close to a *Daily Me*, just for you, in a matter of seconds. In fact that's happening

every day. If the algorithm knows that you like certain kinds of music, it might know, with a high probability, what kinds of movies and books you like, and what political candidates will appeal to you. And if it knows what websites you visit, it might well know what products you're likely to buy, what you think about climate change and immigration.

[6] A small example: Facebook probably knows your political convictions, and it can inform others, including candidates for public office, of what it knows. It categorizes its users as very conservative, conservative, moderate, liberal, and very liberal. It does so by seeing what pages you like. If you like certain opinions but not others, it is easy to put together a political profile. If you mention certain candidates favorably or unfavorably, categorization is easier still. By the way, Facebook doesn't hide what it is doing. On the Ad Preferences page on Facebook, you can look under "Interests," and then under "More," and then under "Lifestyle and culture," and finally under "US Politics," and the categorization will come right up.

[7] Machine learning can be used (and probably is being used) to produce fine-grained distinctions. It is easy to imagine a great deal of sorting—not just from the political right to the political left, but also with specifics about the issues that you care most about, and your likely views on those issues (immigration, national security, equality, and the environment). To say the least, this information can be useful to others—campaign managers, advertisers, fund-raisers and liars, including political extremists. Or consider the hashtag. With #Ireland, #SouthAfrica, #DemocratsAreCommunists, or #ClimateChangelsAHoax, you can find in an instant a large number of items that interest you, or that fit with or even fortify your convictions. The whole idea of the hashtag is to enable people to find tweets and information that interests them. It's a simple and fast sorting mechanism. You can create not merely a *Daily Me* but rather a *MeThisHour* or a *MeNow*. (#*MeNow*? I thought I just made that up, but of course it's in common use.) Many people act as hashtag entrepreneurs; they create or spread hashtags as a way of promoting ideas, perspectives, products, persons, supposed facts, and eventually actions.

[8] Many of us are applauding these developments, which can obviously increase fun, convenience, learning, and entertainment. Al-most no one wants to see advertisements for products that don't interest them. If they're bored by stories about France's economy, why should they have to see such stories on their computer screen or their phone? It is a fair question, but the architecture of control has a serious downside, raising fundamental questions about freedom, democracy, and self-government. What are the social preconditions for a well-functioning system of democratic deliberation or individual liberty itself? Might *serendipity* be important, even if people do not want it? Might a perfectly controlled communications universe—a personalized feed—be its own kind of dystopia? How might social media, the explosion of communications options, machine learning, and artificial intelligence alter the capacity of citizens to govern themselves?

[9] As we will see, these questions are closely related. My largest plea here, in fact, is for *an architecture of serendipity*—for the sake of individual lives, group behavior, innovation, and democracy itself. *To the extent that social media allow us to create our very own feeds, and essentially live in them, they create serious problems.* And to the extent that providers are able to create something like personalized experiences or gated communities for each of us, or our favorite topics and preferred groups, we should be wary. Self-insulation and personalization are solutions to some genuine problems, but they also spread falsehoods, and promote polarization and fragmentation. An architecture of serendipity counteracts homophily, and promotes both self-government and individual liberty.

[10] There is an important clarification. These are claims about the nature of freedom, personal and political, and the kind of communications system that best serves a democratic order. These are not claims about what all or most people are doing. As we will see, many people do like echo chambers, and they very much want to live in them. Many other people dislike echo chambers; they are curious, even intensely so, and they want to learn about all sorts of topics and many points of view. Many people simply gravitate, by default, to the most well-known or popular sites, which do not have a clear ideological orientation. Empirical work confirms these claims, showing that many members of the public are keenly interested in seeing perspectives that diverge from their own, and also that with online browsing, most people spend their time on mainstream sites lacking identifiable political convictions. Many people are open-minded, and shift their views on the basis of what they learn. Such people have an identifiable civic virtue; they are not too sure that they are right, and they want to discover the truth. Many other people much prefer to hear opinions that are consistent with their own, but they are also perfectly willing to hear opinions that challenge them; they do not love the idea of an echo chamber, and they do not create one for themselves. In due course, I will have a fair bit to say about how people are actually using web-sites and social media, and the extent to which people are moving toward an architecture of control. But my central claims are not empirical; they are about individual and social ideals. They are about the kind of culture that is best suited to a well-functioning democracy.

TWO REQUIREMENTS

[11] What I will be emphasizing, then, is people's growing power to filter what they see, and also providers' growing power to filter for each of us, based on what they know about us. In the process of discussing these powers, I will attempt to provide a better understanding of the meaning of freedom of speech in a self-governing society. A large part of my aim is to explore what makes for a well-functioning system of free expression. Above all, I urge that in a diverse society, such a system requires far more than restraints on government censorship and respect for individual choices. For the last several decades, this has been the preoccupation of American law and politics, and in fact the law and politics of many other nations as well, including, for example, Germany, France, England, Italy, South Africa, and Israel. Censorship is indeed the largest threat to democracy and freedom. But an exclusive focus on government censorship produces serious blind spots. In particular, a well-functioning system of free expression must meet two distinctive requirements.

[12] **First**, people should be exposed to materials that they would not have chosen in advance. Unplanned, unanticipated encounters are central to democracy itself. Such encounters often involve topics and points of view that people have not sought out and perhaps find quite irritating— but that might nevertheless change their lives in fundamental ways. They are important to ensure against fragmentation, polarization, and extremism, which are predictable outcomes of any situation in which like-minded people speak only with themselves. In any case, truth matters. I do not suggest that government should force people to see things that they wish to avoid. But I do contend that in a democracy deserving the name, lives—including digital ones— should be structured so that people frequently come across views and topics that they have not specifically selected. That kind of structuring is, in fact, a form of choice architecture from which individuals and groups greatly benefit. Here, then, is my plea for serendipity.

[13] **Second**, many or most citizens should have a wide range of common experiences. Without shared experiences, a heterogeneous society will have a much more difficult time addressing social problems. People may even find it hard to understand one another. Common experiences, emphatically including the common experiences made possible by social media, provide a form of social glue. A national holiday is a shared experience. So is a major sports event (the Olympics or the World Cup), or a movie that transcends individual and group differences (Star Wars is a candidate). So is a celebration of some discovery or achievement. Societies need such things. A system of communications that radically diminishes the number of such experiences will create a range of problems, not least because of the increase in social fragmentation.

[14] As preconditions for a well-functioning democracy, these requirements— chance encounters and shared experiences— hold in any large country. They are especially important in a heterogeneous nation— one that faces an occasional danger of fragmentation. They have even more importance as many nations become increasingly connected with others (Brexit or no Brexit) and each citizen, to a greater or lesser degree, becomes a "citizen of the world." That is a controversial idea, but consider, for example, the risks of terrorism, climate change, and infectious diseases. A sensible perspective on these risks and others like them is impossible to obtain if people sort themselves into echo chambers of their own design. And at a national level, gated communications communities make it extremely difficult to address even the most mundane problems.

[15] An insistence on chance encounters and shared experiences should not be rooted in nostalgia for some supposedly idyllic past. With respect to communications, the past was hardly idyllic. Compared to any other period in human history, we are in the midst of many extraordinary gains, not least from the standpoint of democracy itself. For us, nostalgia is not only unproductive but also senseless. Things are getting better, not worse. Nor should anything here be taken as a reason for "optimism" or "pessimism"— two potential obstacles to clear thinking about new technological developments. If we must choose between them, by all means let us choose optimism.' But in view of the many potential gains and losses inevitably associated with massive technological change, any attitude of optimism or pessimism is far too general to be helpful . . . As we will see, it is much too simple to say that any system of communications is desirable if and because it allows individuals to see and hear what they choose. Increased options are certainly good, and the rise of countless niches has many advantages. But unanticipated, unchosen exposures and shared experiences are important too.

WHY THIS MATTERS: VIOLENCE, PARTYISM, AND FREEDOM

[16] Do echo chambers matter? Exactly why? Some people might not love it if their fellow citizens are living in information cocoons, but in the abstract, that is up to each of us, a reflection of our freedom to choose. If people like to spend their time with Mozart, football, climate change deniers, or Star Wars, so what? Why worry? The most obvious answer is also the narrowest: *violent extremism*. If like-minded people stir one another to greater levels of anger, the consequences can be literally dangerous. Terrorism is, in large part, a problem of

hearts and minds, and violent extremists are entirely aware of that fact. They use social media to recruit people, hoping to increase their numbers or inspire "lone wolves" to engage in murderous acts. They use social media to promote their own view of the world, hoping to expand their reach. The phenomena to be discussed here are contributors to many of the most serious threats we face in the world today.

[17] More broadly, echo chambers create far greater problems for actual governance, even if they do not produce anything like violence or criminality. Most important, they can lead to terrible policies or a dramatically decreased ability to converge on good ones. Suppose (as I believe) that the United States should enact reasonable controls on gun purchases saying, for example, that those on terrorist watch lists should not be allowed to buy guns, unless they can show that they present no danger. Or suppose (as I also believe) that some kind of legislation controlling greenhouse gas emissions would be a good idea. (Perhaps you disagree with these illustrations; if so, choose your own.) In the United States, political polarization on such issues is aggravated by voters' self-segregation into groups of like-minded people, which can make it far more difficult to produce sensible solutions. Even if the self-segregation involves only a small part of the electorate, they can be highly influential, not least because of the intensity of their beliefs. Public officials are accountable to the electorate, and even if they would much like to reach some sort of agreement, they might find that if they do so, they will put their electoral future on the line. Social media certainly did not cause the problem, but in #Republic, things are worse than they would otherwise be . . .

[18] Over the last generation, the United States has seen an explosion in "partyism"—a kind of visceral, automatic dislike of people of the opposing political party. Partyism certainly isn't as horrible as racism; no one is enslaved or turned into a lower caste. But according to some measures, partyism now exceeds racism. In 1960, just 5 percent of Republicans and 4 percent of Democrats said that they would feel "displeased" if their child married outside their political party.⁶ By 2010, those numbers had reached 49 and 33 percent, respectively—far higher than the percentage of people who would be "displeased" if their child married someone with a different skin color.⁶ In hiring decisions, political party matters: many Democrats do not want to hire Republicans, and vice versa, to such an extent that they would favor an inferior candidate of their preferred political party.⁷ Here as elsewhere, we should be cautious before claiming causation; it would be reckless to say that social media and the Internet more generally are responsible for the remarkable increase in partyism. But there is little doubt that a fragmented media market is a significant contributing factor.

[19] By itself, partyism is not the most serious threat to democratic self-government. But if it decreases government's ability to solve serious problems, then it has concrete and potentially catastrophic consequences for people's lives. I have offered the examples of gun control and climate change; consider also immigration reform and even infrastructure—issues on which the United States has been unable to make progress in recent years, in part because of the role of echo chambers. To be sure, the system of checks and balances is designed to promote deliberation and circumspection in government, and prevent insufficiently considered movement. But paralysis was hardly the point—and a fragmented communications system helps to produce paralysis.

[20] There is another problem. Echo chambers can lead people to believe in falsehoods, and it may be difficult or impossible to correct them. Falsehoods take a toll. One illustration is the belief that President Barack Obama was not born in the United States. As falsehoods go, this one is not the most damaging, but it both reflected and contributed to a politics of suspicion, distrust, and sometimes hatred. A more harmful example is the set of falsehoods that helped produce the vote in favor of "Brexit" (the exodus of the United Kingdom from the European Union) in 2016. Even if Brexit was a good idea (and it wasn't), the vote in its favor was made possible, in part, by uses of social media that badly misled the people of the United Kingdom. In the 2016 presidential campaign in the United States, falsehoods spread like wildfire on Facebook. Fake news is everywhere. To date, social media have not helped produce a civil war, but that day will probably come. They have already helped prevent a coup (in Turkey in 2016).

[21] These are points about governance, but, as I have suggested, there is an issue about individual freedom as well. When people have multiple options and the liberty to select among them, they have freedom of choice, and that is exceedingly important. As Milton Friedman emphasized, people should be "free to choose." But freedom requires far more than that. It requires certain background conditions, enabling people to expand their own horizons and to learn what is true. It entails not merely satisfaction of whatever preferences and values people happen to have but also circumstances that are conducive to the free formation of preferences and values. The most obvious way to curtail those circumstances is censorship and authoritarianism—the boot on the face, captured by George Orwell's *Nineteen Eighty-Four*: "If you want a vision of the future, imagine a boot stamping on a human face—forever." A world of limitless choices is incalculably better than that. But if people are sorting themselves into communities of like-minded types, their own freedom is at risk. They are living in a prison of their own design.

DEATH AND LIFE

[22] Let me now disclose a central inspiration for this book, one that might seem far afield: *The Death and Life of Great American Cities* by Jane Jacobs.' Among many other things, Jacobs offers an elaborate tribute to the sheer diversity of cities— to public spaces in which visitors encounter a range of people and practices that they could have barely imagined, and that they could not possibly have chosen in advance. As Jacobs describes great cities, they teem and pulsate with life:

It is possible to be on excellent sidewalk terms with people who are very different from oneself and even, as time passes, on familiar public terms with them. Such relationships can, and do, endure for many years, for decades.... The tolerance, the room for great differences among neighbors— differences that often go far deeper than differences in color— which are possible and normal in intensely urban life . . . are possible and normal only when streets of great cities have built-in equipment allowing strangers to dwell in peace together. . . . Lowly, unpurposeful and random as they may appear, sidewalk contacts are the small change from which a city's wealth of public life may grow.

[23] Jacobs's book is about architecture, not communications. But with extraordinary vividness, Jacobs helps show, through an examination of city architecture, why we should be concerned about a situation in which people are able to create communications universes of their own liking. Her "sidewalk contacts" need not occur only on sidewalks. The idea of "architecture" should be taken broadly rather than narrowly. Websites have architectures, and so do Facebook, Twitter, YouTube, and Reddit. And acknowledging the benefits that Jacobs finds on sidewalks, we might seek to discover those benefits in many other places. At its best, I believe, a system of communications can be for many of us a close cousin or counterpart to a great urban center (while also being a lot safer, more convenient, and quieter). For a healthy democracy, shared public spaces, online or not, are a lot better than echo chambers.

[24] In a system with robust public forums, such as streets and parks, and general-interest intermediaries, such as daily newspapers and network television, self-insulation is more difficult; echo chambers are much harder to create; and people will frequently come across views and materials that they would not have chosen in advance. For diverse citizens, this provides something like a common framework for social experience. "Real-world interactions often force us to deal with diversity, whereas the virtual world may be more homogeneous, not in demographic terms, but in terms of interest and outlook. Place-based communities may be supplanted by interest-based communities." Consider here the finding that communities that believed the apocalypse was near, and thought the attacks on September 11, 2001, were a clear sign to that effect, used the Internet so as "to insulate" themselves "from the necessarily divergent ideas that might generate more constructive public discussion." To be sure, we do not yet know whether anything can or should be done about fragmentation and excessive self-insulation. I will take up that topic in due course. For purposes of obtaining understanding, few things are more important than to separate the question of whether there is a problem from the question of whether anything should be done about it. Dangers that cannot be alleviated continue to be dangers. They do not go away if or because we cannot, now or ever, think of decent solutions. It is much easier to think clearly when we appreciate that fact.

WHAT FACEBOOK WANTS

[25] On June 29, 2016, Facebook made a significant announcement, under a post called "Building a Better News Feed for You..." It didn't exactly say that it had found a way to produce a *Daily Me*, but it came fairly close, and it made clear its aspirations. The post emphasizes that "the goal of News Feed is to show people the stories that are most relevant to them." With that point in mind, why does Facebook rank stories in its News Feed? "So that people can see what they care about first, and don't miss important stuff from their friends." In fact, the News Feed is animated by "core values," starting with "getting people the stories that matter to them most." Facebook therefore asks this question: "If you could look through thousands of stories every day and choose the ten that were most important to you, which would they be? The answer should be your News Feed. It is subjective, personal, and unique— and defines the spirit of what we hope to achieve." (It's worth pausing over that.) I should note that I like Facebook and use it regularly— but it can improve.

[26] Consistent with that spirit, Facebook says, "To help make sure you don't miss the friends and family posts you are likely to care about, we put those posts toward the top of your News Feed. We learn from you and adapt over time. For example, if you tend to like photos from your sister, we'll start putting her posts closer to the top of your feed so you won't miss what she posted while you were away." In this way and others, personalization matters: "Something that one person finds informative or interesting may be different from what another person finds informative or interesting." The News Feed is designed so that different people get what they want.

[27] Facebook says that it does not play favorites. Its business is "connecting people and ideas— and matching people with the stories they find most meaningful." (The word "meaningful" is interesting here. What does it mean?) It follows that "as News Feed evolves, we'll continue building easy-to-use and powerful tools to give you the most personalized experience." That "we" is unduly confident that "the most personalized experience" is what is most desirable. From the post, it is not exactly clear what Facebook did to improve the situation, but the company appears to have altered its algorithm to ensure that at the top of your News Feed, you will see items from your friends, thus increasing the likelihood that what you will see will be what most interests you. The post concludes: "We view our work as only 1% finished— and are dedicated to improving along the way." That's good news.

[28] We do not know for sure, but Facebook probably made this change for three reasons. First, it had recently faced allegations of political bias, in the form of suppression of conservative news sources. An algorithm that emphasizes family and friends, and seemingly puts users in full control, can claim political neutrality. Second, Facebook has an obligation to its shareholders, and if its News Feed really can be turned into a *Daily Me*, it might well get more clicks, which means more revenue. Third, many users had been merely posting news articles of various sorts, which meant a reduction in original posts. People might find the reposted articles less interesting, and if so, there are fewer clicks, making for a less attractive product. (I speculate that the third reason might be the most important.)

[29] It is entirely reasonable for Facebook to take these points into account. But we should not aspire to a situation in which everyone's News Feed is perfectly personalized, so that supporters of different politicians— Bernie Sanders, Hillary Clinton, Donald Trump, someone else— see fundamentally different stories, focusing on different topics or covering the same topics in radically different ways. Facebook seems to think that it would be liberating if every person's News Feed could be personalized so that people see only and exactly what they want. Don't believe it. In the 2016 presidential campaign, the News Feed spread a lot of falsehoods. Facebook is right to underscore the importance of core values, but it might want to rethink its own. True, it is a business, not a public utility. True, it has obligations to its shareholders. But in view of its massive role in determining what kinds of news people see, it is far from ideal if it does not include, among its core values, promoting or at least not undermining democratic self-government. Facebook can do better.

"I'M SCARED; HOW ABOUT YOU?"

[30] I have a friend who has a rule: "You cannot be happier than your spouse." That might be too simple, but he's onto something: emotions are contagious. If you are in a happy family, you'll be happier yourself, and if your partner or your children are enraged about something, or frightened, your own emotions will tend in the same direction. It stands to reason that the emotional valence of what you read or see will have analogous effects. If your Twitter feed is full of pessimistic people, verging on despair about the economy or the fate of your nation, you'll become more pessimistic as well. A more alarming possibility: if an alienated young person is reading material from a terrorist organization, furious about the supposed misdeeds of the United States or the United Kingdom, he might get furious too— and perhaps be led to commit acts of violence. One consequence of personalization is likely to be not only fragmentation with respect to topics and points of view but also fragmented *feelings*— perhaps in general, or perhaps with respect to specific objects and positions.

[31] Evidence to this effect comes from an important and controversial study by Facebook itself. In the study, Facebook worked with Cornell University to conduct an experiment in which the company deliberately fed certain users sad posts in order to test whether the sadness of those posts would affect the emotions of those users. Of course Facebook did not have direct access to users' emotions— but it could see what they did next. Would their own behavior be affected? Would their posts shift in some way? Yes and yes. As it turned out, the users who were given the sad posts began posting sad posts themselves. If we measure the effects on their emotions by what they did next, we can fairly say that sadness proved contagious on Facebook pages, just as in families and workplaces. The study is controversial because Facebook users do not exactly love the idea that the company may be manipulating their feelings. (If Facebook really wanted to make its users, or some subcategories of them, mad or sad, it could easily do that.) But to its credit, Facebook made a genuine contribution to science, producing as it did strong evidence that the emotional valence of what you read on social media will affect not only what you think but also how you feel. And if people are sorting themselves into different groups, or being sorted into such groups, it is inevitable that the emotional experiences of those groups will differ—often in response to precisely the same events.

[32] A World Cup game is a benign example; if Germany is playing Argentina, fans of different teams will have different emotional reactions to the same outcome. In many ways, an election is similar. But we could also have radically different emotional responses to an event that is not self-evidently polarizing— a terrorist attack, a natural disaster, or a purely scientific report. In fact that is happening every day, in large part because of the power of echo chambers.