

#Rep
Sunstein
kind of
cieties

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.

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.

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, and 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

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.² 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, and what you think about climate change and immigration.

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.

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, fundraisers and liars, including political extremists.

Or consider the hashtag. With #Ireland, #SouthAfrica, #DemocratsAreCommunists, or #ClimateChangeIsAHoax, 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.

Many of us are applauding these developments, which can obviously increase fun, convenience, learning, and entertainment. Almost 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?

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.

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.³ Many people are open-minded, and their views shift 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 websites 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

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.

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.

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.

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.

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.⁴ 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. Automobiles are great, but in the United States alone, many thousands of people die every year in car crashes. Plastics are a huge advance, but they have created a serious waste disposal problem. What I mean to provide is not a basis for pessimism but instead a lens through which we might understand, a bit better than before, what makes a system of freedom of expression successful in the first place, and what a well-functioning democracy requires. That improved understanding will equip us to understand a free nation’s own aspirations, and thus help us to evaluate continuing changes in the system of communications. It will also point the way toward a clearer understanding of the nature of citizenship and its cultural prerequisites.

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

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.

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.

I have worked in various capacities with the federal government and met on many occasions with members of Congress. With respect to important issues, Republicans have said to me, "Of course we would like to vote with the Democrats on that one, but if we did, we would lose our jobs." There is no question that behind closed doors, Democrats would on occasion say the same thing about working with Republicans. Both sides are worried about the effects of echo chambers—about an outburst of noisy negativity from segments of constituents, potentially producing serious electoral retribution. Social media increase the volume of that noise, and to that extent, they heighten polarization.

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.⁵ 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.

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.

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).

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

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.⁸ 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.⁹

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.

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.”¹⁰ 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.”¹¹

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

On June 29, 2016, Facebook made a significant announcement, under a post called “Building a Better News Feed for You.”¹² 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 10 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.

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.

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 percent finished—and are dedicated to improving along the way.” That’s good news.

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.)

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?"

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.

Evidence to this effect comes from an important and controversial study by Facebook itself.¹³ 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.

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.

PRECURSORS AND INTERMEDIARIES

To some people, unlimited filtering may seem quite strange—a potential product of recent technologies and perhaps even the stuff of science fiction (come true, as it frequently does). But in many ways, it is continuous with what has come before. Filtering is inevitable, a fact of life. It is as old as humanity itself. It is built into our minds. No human being can see, hear, or read everything. In the course of any hour, let alone any day, every one of us engages in massive filtering, simply in order to make life manageable and coherent. Attention is a scarce commodity, and people manage their own attention in order to ensure that they are not overwhelmed.

Indeed, the entire field of behavioral science can be seen to stem from an insistent focus on the limited nature of attention, and the filters we impose on our thought and experience. Daniel Kahneman, Nobel Prize winner and a founder of the field, is widely known for his 2011 masterpiece, *Thinking, Fast and Slow*. But the

arc of his career was nicely signaled by the title of his first book, published in 1973: *Attention and Effort*. Much of behavioral science emphasizes that it is effortful to attend to certain topics and concerns. People often want to minimize that effort. That's built into our species. Sometimes we allocate our attention deliberately: we decide to focus on our children, not on problems in Syria and Iraq. But often we allocate our attention without thinking about it. When you're driving, you concentrate on what's in front of you and what's in back, and many of your motions are automatic. What we "see" and what we notice are frequently outside our conscious control.

With respect to the world of communications, a free society gives people a great deal of power to filter out unwanted materials. Only tyrannies force people to read or watch. In free nations, those who read newspapers do not read the same newspaper; many people do not read any newspaper at all. Every day, people make choices among magazines based on their tastes and point of view. Sports enthusiasts choose sports magazines, and in many nations they select a magazine focused on the sport of their choice—*Basketball Weekly*, say, or the *Practical Horseman*. Conservatives can read *National Review* or the *Weekly Standard*; countless magazines are available for those who like cars; *Dog Fancy* is a popular item for canine enthusiasts; people whose political views are somewhat left of center might like the *American Prospect*; many people like *Cigar Aficionado*.

These are simply contemporary illustrations of a long-standing fact of life in many countries: a diversity of communications options and a wide range of possible choices. But the emerging situation does contain large differences, stemming above all from dramatic increases in individual control over content, the number of available options, the sheer speed with which people can receive information, and corresponding decreases in the power of *general-interest intermediaries*.¹⁴

General-interest intermediaries include newspapers, magazines, and broadcasters. An appreciation of their social functions will play a large role in this book. As prominent current examples, consider *Time*, *Newsweek*, the *New York Times*, the *Wall Street Journal*, the

Columbia Broadcasting System, and the *New York Review of Books*. People who rely on such intermediaries have a range of chance encounters, involving shared experiences with diverse others, and also exposure to materials and topics that they did not seek out in advance. The *New York Review of Books* provides you with a lot of material that you would not have chosen in advance; so too with the daily newspaper. You might find a range of stories that you would not have selected if you had the power to include or exclude them. Your eyes might come across a story about tensions over immigration in Germany, crime in Los Angeles, innovative business practices in Tokyo, a terrorist attack in India, or a hurricane in New Orleans, and you might read those stories, although you would hardly have placed them in your Twitter feed or your Daily Me. You might watch a particular television channel—perhaps you prefer channel 4—and when your favorite program ends, you might see the beginning of another show, perhaps a drama or news special that you would not have chosen in advance, but that somehow catches your eye.

Reading *Time* or *Newsweek*, you might come across a discussion of endangered species in Madagascar or genocide in Darfur, and it might interest you, even affect your behavior, maybe change your life, despite the fact that you would not have sought it out in the first instance. A system in which individuals lack control over the particular content that they see has a great deal in common with a public street, where you might encounter not only friends but also a heterogeneous array of people engaged in a wide array of activities (including perhaps bank presidents, political protesters, and panhandlers).

Some people believe that the mass media are dying—that the whole idea of general-interest intermediaries, providing shared experiences and exposure to diverse topics and ideas for millions, was a short episode in the history of human communications. As a prediction, this view seems wrong; even on the Internet, the mass media continue to play a large role. But certainly their significance has been falling over time.

It is an understatement to say that the communications market is in flux. Many of the most important general-interest intermediaries

are in serious trouble. We should not forget that from the standpoint of human history, even in industrialized societies, such intermediaries are relatively new and far from inevitable. Newspapers, radio stations, and television broadcasters have particular histories with distinctive beginnings as well as possibly distinctive endings. In fact, the twentieth century should be seen as the great era for the general-interest intermediary, which provided similar information and entertainment to millions of people.

The twenty-first century may well be altogether different on this score. Consider one small fact: in 1930, daily newspaper circulation was 1.3 per household—a rate that had fallen to less than 0.50 by as early as 2003, even though the number of years of education, typically correlated with newspaper readership, rose sharply in that period. At the very least, the sheer volume of options and the power to customize are sharply diminishing the social role of the general-interest intermediary.

Indeed, one of the distinguishing features of the current era, accompanying the Daily Me, is the *special-interest intermediary*. Instead of serving as broad sources of information that cover a variety of topics, online news outlets often take the form of specialized “verticals” that focus on narrower subjects, such as sports, technology, or politics, or use specialized methodologies of interest to niche markets, whether large or small (such as *fivethirtyeight.com*, which emphasizes statistical approaches to politics and sports). These outlets are proliferating at a rapid rate; they attract capital from investors and are run more like start-ups than established news outlets. The greater specialization of these information sources, such as the various platforms run by Vox Media, will produce some echo chambers—and to that extent, diminish the likelihood of shared experiences.

HER

For a vivid illustration of what’s around the corner, consider Spike Jonze’s brilliant 2015 film, *Her*. In some ways, I think that it ranks

with *Nineteen Eighty-Four* or *Brave New World* as a depiction of a humanly recognizable dystopian future, and it captures a dystopia that Orwell and Huxley could not have envisioned.

Theodore Twombly, the film’s protagonist, makes a living by writing highly personalized notes and cards—for example, anniversary notes from wives to husbands—based on a great deal of information about both the sender and the receiver. In Twombly’s world, love letters are simultaneously outsourced and customized. Twombly isn’t exactly an operating system, but he sure acts like one. He also faces an imminent divorce, and his own life is in shambles, filled with video games and anonymous phone sex (personalized, of course). Everything changes when he purchases an operating system, a form of artificial intelligence (think Siri 4.0) who names “herself” Samantha.

Samantha has access to Twombly’s computer, including his e-mails. She is a fast reader: she knows what he likes and dislikes, and she understands his strengths and weaknesses. Perhaps above all, she is interested in him. She listens. She yearns to see the world through his eyes. She is there when he wakes up, and every evening she’s the one who says good night to him and to whom he says good night. She watches him while he sleeps.

If that were all, of course, Twombly’s interest would wane quickly. Unless you are an impossible narcissist, you can’t fall for someone whose only words are “Tell me more!” As she is constructed, Samantha has independent interests and concerns. She likes to write music, she’s playful, she’s curious, she can be insecure, and she’s a tease. We can’t know for sure, but perhaps those characteristics are a product of personalization as well. Perhaps they are exactly what Twombly wants and needs. Perhaps the algorithm knows, from a perusal of his browsing habits, how much independence he wants his partner to have, and exactly what kind.

Twombly falls in love with Samantha. How could he avoid that? She knows everything about him. She’s his Daily Me, turned into a lover. Maybe that’s irresistible. Let’s hope that in reality, our operating systems will never become our lovers. (It’s not fanciful to predict

that they will, at least in some sense, though the movie makes a compelling argument for real human beings, imperfect personalization and all.) Whatever happens, we can take *Her* as a metaphor for processes, occurring every minute, by which our browsing habits—the words we use, the places we go, the friends we make, the things that we “like”—provide countless clues about our tastes and values. You don’t have to be Samantha to be able to cater to those tastes and values. You can be some algorithm now in common use.

SOCIAL MEDIA

My topic is online behavior in general, and so much of the discussion will involve uses of big websites—*newyorktimes.com*, *foxnews.com*, *Amazon.com*, *pandora.com*. But I will also spend a lot of time on social media, which requires some definitional work.

Speaking of pornography, Supreme Court Justice Potter Stewart famously wrote, “I know it when I see it.” Do we know social media when we see it? Any particular examples will become dated, but Facebook, Twitter, Instagram, and Snapchat certainly count. According to a helpful definition, social media are “Internet-based platforms that allow the creation and exchange of user-generated content, usually using either mobile or web-based technologies.”¹⁵ Wikipedia fits that definition, because people use it to produce content. YouTube must be included, because people share content there; Flickr and Vine are also examples. Blogs (such as *Marginal Revolution*) and microblogs (such as Twitter) are definitely included. So are social networking sites, most prominently Facebook, but also WhatsApp, Orkut, Yik Yak, Tumbler, and Tuenti. Social media can be used both for social purposes and games (such as *Second Life* and *Pokémon Go*). Can apps count? For my purposes, they certainly can, so long as they fit the definition.

It should be clear that we are dealing with a highly protean category, and its content changes rapidly over time. In 2006, blogs and the blogosphere were all the rage; while blogs exist and remain important, they have far less centrality (and the very word

“blogosphere” seems to be a relic, a bit like “rotary phone” or “groovy”). Twitter was launched in 2006, Tumbler and WhatsApp in 2010, and Snapchat in 2011. Social media often have nothing at all to do with politics or democracy (indeed, they are a kind of vacation from it), and to that extent, they do not trigger my principal concerns here. But even if they are wholly apolitical, they might create niches, and niches produce fragmentation.

A WORD ON BASELINES

Any assessment of the world of the Internet, and any claims about what’s wrong with it, must ask one question: *Compared to what?* We could easily imagine a suggestion that in some prior period—say, 1940, 1965, or 1980—the world of communications was much better. Perhaps there was a golden age of communication; many people think so. But that is not my claim here. I will be comparing the current situation not to some lost utopia but instead to a communications system that has never existed—one in which existing technological capacities and unimaginable improvements are enlisted to provide people with the equivalent of a great city, full of substance, fun, diversity, challenge, comfort, disturbance, colors, and surprise.

That is frustratingly vague, I know. It might help to say what the idealized baseline does *not* include. (It is much easier to speak of injustices than to offer an account of justice.) It does not involve a system of acute political polarization, in which large numbers of people sort themselves into information cocoons. It is not highly fragmented. It involves unanticipated exposures to topics and ideas. It counteracts falsehoods, spread by innocent or not-so-innocent people, misleading their fellow citizens about issues of health and wealth. It promotes deliberation among people who are not of like mind. It recognizes that some people are curious, and it cultivates political curiosity, seeing it as a civic virtue. (Recall that identifiable people have that virtue and like to read material that challenges their own preconceptions.)

These ideas are vague too. We should not offer a conception of an ideal communications market that is sanctimonious or preachy, or hopelessly ill suited to the lives (not to mention the attention spans) of actual human beings. My hope is that the ideal baseline, and departures from it, will emerge as we explore concrete problems.

POLITICS, FREEDOM, AND FILTERING

In the course of the discussion, we will encounter many issues. Each will be treated in some detail, but for the sake of convenience, here is a quick catalog:

- the importance of chance encounters and shared experiences for democratic societies
- the large difference between pure populism, or direct democracy, and a democratic system that attempts to ensure deliberation and reflection as well as accountability
- the intimate relationship between free speech rights and social well-being, which such rights often serve
- the pervasive risk that discussion among like-minded people will breed excessive confidence, extremism, contempt for others, and sometimes even violence
- the potentially dangerous role of social cascades, including “cybercascades,” in which information, whether true or false, spreads like wildfire
- the enormous potential of the Internet and other communications technologies for promoting freedom in both poor and rich countries
- the utterly implausible nature of the view that free speech is an “absolute”
- the ways in which information provided to any one of us is likely to benefit many of us

- the difference between our role as citizens and our role as consumers
- the inevitability of regulation of speech, and indeed the inevitability of speech regulation benefiting those who most claim to be opposed to “regulation”
- the potentially destructive effects of intense market pressures on both culture and government

But the unifying issue throughout will be the various problems for a democratic society that might be created by the power to filter. Democracies may or may not be fragile, but polarization can be a serious problem, and it is heightened if people live in different communications universes—as in fact they sometimes seem to do in the United States, the United Kingdom, France, Germany, and elsewhere. There is no doubt that the modern communications environment, including social media, contributes to the rise of partyism.

One question, which I answer in the affirmative, is whether individual choices, innocuous and perfectly reasonable in themselves, might produce a large set of social difficulties. Another question, which I also answer in the affirmative, is whether it is important to maintain the equivalent of “street corners” or “commons” where people are exposed to things quite involuntarily. More specifically, I seek to defend a particular conception of democracy—a deliberative conception—and evaluate, in its terms, the outcome of a system with perfect power of filtering.

My claim is emphatically not that street corners and general-interest intermediaries will or would disappear in a world of perfect filtering. To what degree the market will produce them or their equivalent is an empirical issue. Some people invite general-interest intermediaries by default; if they are looking for news, that is where they go, and they do not much care about ideological disposition. Some people have a strong taste for street corners and their equivalent on television and the Internet. Indeed, the Internet

holds out immense promise for allowing people to be exposed to materials that used to be too hard to find, including new topics and new points of view. If you would like to find out about different forms of cancer and different views about possible treatments, you can do so in less than a minute. If you are interested in learning about the safety record of different automobiles, a quick search will tell you a great deal. If you would like to know about a particular foreign country, from its customs to its politics to its weather, you can do better with the Internet than you could have done with the best of encyclopedias.

From the standpoint of those concerned with ensuring access to more topics and more opinions, existing communications technologies are a terrific boon. But it remains true that many apparent “street corners,” on the Internet in particular, are highly specialized, limited as they are to particular topics and points of views. What I will argue is not that people lack curiosity or that street corners will disappear but instead that there is an insistent need for them, and that a system of freedom of expression should be viewed partly in light of that need. In particular, I will emphasize the risks posed by any situation in which hundreds of thousands, millions, or even hundreds of millions of people are mainly listening to louder echoes of their own voices.

WHAT ISN'T THE ISSUE

Some clarifications, designed to specify the central issues, are now in order. I will be stressing problems on the “demand” side of the speech market. These are problems that stem not from the actions of producers but instead from the choices and preferences of consumers. I am aware that on one view, the most important emerging problems come from large corporations, and not from the many millions and indeed billions of individuals who make communications choices. In the long run, however, I believe that the most interesting questions, and certainly the most neglected ones, involve consumer behavior. This is not because consumers are usually confused,

irrational, or malevolent. It is because choices that seem perfectly reasonable in isolation may, when taken together, end up disserving democratic goals.

Because of my focus on the consumers of information, I will not be discussing a wide range of issues that have engaged attention in recent decades. Many of these issues involve the allegedly excessive power of large corporations or conglomerates.

I will not deal with the feared disappearance of coverage of issues of interest to small or disadvantaged groups. Every day, that is less of a problem. On the contrary, there has been a tremendous growth in niche markets, serving groups both large and small. With a decrease in scarcity, this trend will inevitably continue. In #Republic, people should be able to find what they want, and should be able to become members of groups that they like. Technological development is a great ally of little groups and minorities, however defined. People with unusual or specialized tastes are not likely to be frozen out of the emerging communications universe. The opposite is much more likely to be true: they will have easy access to their preferred fare—far easier than ever before. If you love *Star Wars*, the 2012 television show *Awake*, or Taylor Swift, you can find people who will share the love.

I will not be exploring the fascinating increase in people's ability to participate in *creating* widely available information—through art, movies, books, science, and much more. With social media, any one of us might be able to make a picture, a story, or a video clip available to all of us; YouTube is merely one example. In this way, social media have a powerful democratizing function.¹⁶ Countless websites are now aggregating diverse knowledge. For diverse products—books, movies, cars, doctors, and computers—it is easy to find sources that tell you what most people think, and it is easy as well to contribute to

that collective knowledge. Prediction markets, for example, aggregate the judgments of numerous forecasters, and they are proving to be remarkably accurate. There is much to be said about the growing ability of consumers to be producers too.¹⁷ But that is not my topic here.

- I will provide little discussion of monopolistic behavior by suppliers, or manipulative practices by them. Undoubtedly some suppliers do try to monopolize, and some do try to manipulate; consider, for example, the fact that Google provides paid links for certain sites (but not others) or tailors search algorithms to present certain search results (over others). Every sensible producer of communications knows that a degree of filtering is a fact of life. Producers also know something equally important but less obvious: consumers' attention is the crucial (and scarce) commodity in the emerging market. Companies stand to gain a great deal if they can shift attention in one direction rather than another. This is why many Internet sites supply information and entertainment to consumers for free. Consumers are actually commodities, and they are often "sold" to advertisers in return for money; it is therefore advertisers and not consumers who pay. This is pervasively true of radio and television.¹⁸ It is true of numerous websites too.

Especially in light of the overriding importance of attention, some private companies will attempt to manipulate consumers, and occasionally they will engage in monopolistic practices. Is this a problem? No unqualified answer would make sense. A key question is whether market forces will reduce the adverse effects of efforts at manipulation or monopoly. I believe that to a large extent, they will, because competition for eyeballs is fierce, but that is not entirely clear. For example, Facebook is no ordinary competitor, and it has a lot of market power. But that is

not my main concern here. For a democracy, many of the most serious issues raised by new technologies do not involve manipulation or monopolistic behavior by large companies. By contrast, personalization via algorithm will be a central theme.

- I will put to one side the active debate over the uses of copyright law to limit the dissemination of material on the Internet and elsewhere. This is an exceedingly important debate, to be sure, but one that raises issues very different from those explored in this book.¹⁹ Nor will I explore the sharply contested question, in some ways related, of "net neutrality," designed to level the playing field among communications providers.

- I will not be discussing the "digital divide," at least not as the term is ordinarily understood. People concerned about that problem emphasize the existing inequality in access to new communications technologies—an inequality that divides those with and without access to the Internet. That is indeed an important issue, certainly domestically and even more so internationally, because it threatens to aggravate existing social inequalities, many of them unjust, at the same time that it deprives many millions (perhaps billions) of people of information and opportunities. But in both the domestic and international context, that problem seems likely to diminish over time, as new technologies, above all the Internet, are made increasingly available to people regardless of their income or wealth.

Of course we should do whatever we reasonably can to accelerate the process, which will provide benefits, not least for both freedom and health, for millions and even billions. But what I will describe will operate even if everyone is on the right side of that divide—that is, even

if everyone has access to the Internet. My focus will be on the distinctive cultural and political divides, across values and tastes, that are emerging in the presence of universal access—on how reasonable choices by individual consumers might produce both individual and social harm. This point is emphatically connected with inequalities, but not in access to technologies; it does not depend in any way on inequalities there.

The digital divides that I will explore may or may not be a nightmare. But if I am right, there is all the reason in the world to reject the view that free markets, as embodied in the notion of “consumer sovereignty,” exhaust the concerns of those who seek to evaluate any system of communications. The imagined world of innumerable, diverse editions of the Daily Me is not a utopian dream, and it would create—is creating—serious problems from the democratic point of view.

AN ANALOGY AND AN IDEAL

The changes now being produced by contemporary communications technologies are understated, not overstated, by the idea of the Daily Me. What is happening goes far beyond the increasingly customized computer screen.

Many of us telecommute instead of going to a workplace; this is a growing trend. Rather than visiting the local bookstore, where we might well end up seeing a number of diverse people, we shop for books on Amazon.com. Others increasingly avoid local restaurants, because seamless.com, or something like it, is entirely delighted to deliver sushi or a pizza to us. Near the dawn of the modern era, media analyst Ken Auletta enthused, “I can sample music on my computer, then click and order. I don’t have to go to a store. I don’t have to get in a car. I don’t have to move. God, that’s heaven.”¹

Really? Heaven? True, if you are interested in anything at all—from computers to linens to diamonds to cars to medical advice—an online company will be happy to assist you. Indeed, if you would like to attend college or even get a graduate degree, you may be able to avoid the campus. College education is available online, and if you’d like to perform marriage ceremonies, you can get licensed to do that too.²

It would be foolish to claim that this is bad, or a loss, in general or on balance. On the contrary, the dramatic increase in convenience is a wonderful blessing for people. Driving around in search of gifts can be a real bother. (Can you remember what this used to be like? Is it still like that for you?) For many of us, the chance to point and click is an extraordinary improvement. And many people, both rich and poor, take advantage of current technologies