3

The Attention Machine

On July 15, 2015, professional golfer Greg Norman appeared on the American morning news show *Today* and challenged presenter Matt Lauer to sit while colleagues dumped a bucket of ice water on his head. Lauer promised to donate money to a hospice service in Florida's Palm Beach County. Lauer then challenged three friends, including news presenter Brian Williams, lifestyle guru Martha Stewart, and radio personality Howard Stern, to sit for the same bracing experience and pledge to donate to the hospice. At that moment, a fundraising gimmick that had been circulating among professional golfers for months erupted into a craze. Within weeks, all over Facebook and YouTube hundreds of thousands of people posted videos of being drenched in ice water for a cause and then challenging three friends to do the same. The challenge was odd: either the person who had been challenged got soaked in ice water or she or he pledged to donate. Often the subject did both. But just as often, the video of the drenching offered no clue about the destination of the charitable contribution. Many videos never mentioned a cause at all.[1]

Soon after Lauer issued his challenge on American television (which meant, of course, that the video of that event would circulate on Facebook as well), word spread that the Ice Bucket Challenge should support the Amyotrophic Lateral Sclerosis (ALS) Association. Soon all the other charities that had been cited by those who committed the act faded into the background. Facebook made it easy for the challenge to spread, because a person could both name those challenged and "tag" them through the service Facebook provides to link profiles to content. Ultimately more than 2.5 million Ice Bucket Challenge videos circulated on Facebook. More than twenty-eight million people engaged with those videos in some way. Celebrities and politicians joined in. And the trend spread far beyond the United States. In just one month in the summer of 2014 the Amyotrophic Lateral Sclerosis Association gathered $98.2 million, compared to just $2.7 million during the same period the year before. In India, where clean water is scarce and at a premium, the "rice bucket challenge" spread, urging people to donate rice to poor people. And in Gaza, activists posted videos of the "rubble bucket challenge" to raise awareness of the misery of daily life in the severed territory that had seen its infrastructure destroyed during conflagrations with Israel.[2]

By simple measures the Ice Bucket Challenge was a spectacular success. The money it raised contributed to essential research against ALS, a rare condition in which parts of the nervous system malfunction, causing progressive weakness and muscle wasting. As the disease progresses, it hinders basic activities such as walking, speaking, swallowing, and breathing. Life expectancy for about half of those with ALS is three to four years from start of diagnosis, but some can live much longer. The effects of ALS are so powerful that generating compassion for those afflicted was not difficult. The challenge has always been to raise awareness of a disease that affects so few people.[3]

There is a problem with all this success, however. In 2014 a larger amount of money went to ALS research than ever before. Was 2014 a particularly crucial year for ALS research? Was the incidence of the disease rising? No, it was like any other year. ALS, like every other debilitating disease, deserves more money for research and treatment. But a dollar donated to ALS research is a dollar denied to malaria prevention, breast cancer research,

HIV treatment and prevention, or the most deadly affliction of all, heart disease. Who should decide which diseases get more research money and which should get less in any given year? On what should we base that judgment? Perhaps we should consider mortality. Perhaps we should consider the numbers of people afflicted. Perhaps we should consider the likelihood of a breakthrough treatment. Many countries employ councils of physicians, scientists, and epidemiologists to make such judgments about the allocation of public research funds.

But as public funding shrinks in many places, including the United States, private sources of research funding become more important. For private sources, whether in large sums from foundations or wealthy donors or in small amounts from millions of Facebook users, publicity makes all the difference. The associations devoted to funding research and treatment rely on narratives of those who suffer. They offer glimmers of hope that the next drive might push the research to its breakthrough point. And since the Ice Bucket Challenge, they try to attach their causes to spectacles that would spread around Facebook.[4]

The cause that makes the catchiest, cutest, or most clever Facebook campaign generates the most financial support. This is a terrible way to determine where charitable contributions should go, especially when medical science is involved. Nothing in any of the Ice Bucket Challenge videos revealed how well the Amyotrophic Lateral Sclerosis Association spends the money it raises. Few if any of the cold and shivering subjects of the videos, for all their good intentions, explained why ALS was such a debilitating and fatal disease or the good to which the money donated would be spent. Many of the videos omitted any reference to ALS. They became spectacles for the sake of spectacle. In the summer of 2014 the world did not discuss and debate whether ALS was the most urgent or even the most promising of causes. It was just the most fun to think about—even though the thinking remained at the level of watching a person get soaked in cold water. The success of the Ice Bucket Challenge pushed all such philanthropic endeavors to consider more lively ways to engage interest. For many, that meant devoting more attention to Facebook, where attention is the only currency that matters.[5]

THE ATTENTION ECONOMY

Attention is scarce. It is limited by time, of course. Most people find it difficult to pay attention to more than one thing at a time. Yet it's easy to steal. A flash, a movement, or an odd sound can pull attention away from this page to something in the periphery of your field of vision. It takes an active decision to move not only your eyes (because they can glaze over when staring at these letters) but your mental processes back to this book. Attention feeds thought. And thought works in streams. If those streams are limited by duration or not allowed to stay steady and focused, the power of that thought diminishes. Yet attention is also valuable, even—or perhaps especially—when it is brief and shallow. People are more easily tickled or angered when attention is fleeting. They are also more likely to be convinced to click on a link on a web page, application, or email if their attention is shallow and short. To sell something, an advertiser must grab your attention, if only for a moment. Advertisers know that while they have your attention, someone else is trying to steal it away from them. So as our media ecosystem becomes more polluted and fractured, each player in it experiments with new designs, targeting strategies, and stimuli to steal attention and then hold it long enough to convince the potential customer to take some action. Add up all those players seeking smaller and smaller fractions of our attention over long periods and the logic of the market generates cacophony. Few moments are spared from the frantic attempts at attention theft. Few spaces remain unmarked by some effort to pull our eyes from the direction we choose.[6]

As many parts of the world deindustrialized over the past forty years, leaving labor with lower wages and fewer protections, we have been promised that the "information economy" would replace industrialization. Information is not scarce. It's too abundant, in fact. Managing and filtering information has become a valuable function as a result. Google and to a lesser extent Facebook help us manage the torrent of information around us by doing the work of deciding what's valuable or interesting to us. Part of that transition to an information economy involved the rise of industries completely devoted to capturing our attention. Again, Google

and Facebook have cornered the market on that. Monetizing our captured attention pays for the labor and technology that enable Google and Facebook to filter the flood of information so effectively. And while those two companies are far from the only players in the attention economy, they are the best at it. Google and Facebook are the most successful of the set of firms that connect advertisers to potential customers. That set includes newspapers, websites, television networks and stations, radio stations, and billboard companies. It now also includes coach and bus services, universities, sports stadia, and hotels, all of which have developed methods of hosting advertising for the patrons and customers who spend time in them. Law professor Tim Wu, who has charted the history of the attention economy, calls this set of firms "attention brokers."[7]

The attention economy emerged from an ideological decision that whole industries made in the late 1990s, when the World Wide Web made it possible to imagine that people would read text and watch video on small LCD screens instead of on printed pages and in dark theaters. Content had to be free. The word "free" carried a lot of weight in those days, signifying the spirit of the age. Capitalism and democracy were ascendant. Libertarianism was no longer considered a fringe ideology. "Free" meant digital content could be cost-free and able to circulate unfettered. To lift the world to new heights of liberty and creativity, the marginal cost of goods and services should approach zero, the theory went. To pay the coders, the utility companies, and the lawyers, however, digital companies would need revenue. Starting with Yahoo in the late 1990s and culminating with Google by the early 2000s, advertising would supply the capital. Once this entire industry decided that content had to be free, it put pressure on industries that had long relied on unit sales or the price of admission, such as music labels, book publishers, and film studios, to drop their prices close to free. These other companies started looking for other sources of revenue that embedded themselves in the attention economy, thus competing against each other not only for the marginal cash in a family's entertainment budget but also for their scarce free time. Culture and entertainment figures strove to make themselves "lifestyle brands" and pushed their images out on YouTube, Twitter, and Instagram. The attention economy has its macroeconomic

phenomena, as exemplified by Google, Facebook, and the entertainment and journalism industries. And it has its microeconomic forces, acting upon everyone else who looked at work life as a matter of free agency and a constant hustle for fleeting opportunities.

The ideology that underlies the attention economy also structures individual experience and cultural expectation. When successful private sector managers meet my students, they often tell these aspiring media workers that they must behave as "brands," constantly and exhaustively promoting themselves through Facebook, Twitter, YouTube, and Instagram. Young people are signaled that becoming a "microcelebrity" may lead to bigger and better macrocelebrity later on. The virtues of not being famous rarely have a champion or a fairy tale behind them. The pinnacle of the struggle for attention, which we are promised will surely pay off through wealth and fame, is the TED Talk. Purposely informal and limited to eighteen minutes, these punchy, pithy talks are meant to inspire and entertain. They don't invite deliberation or debate. They don't demand immersion or even background reading. They are capsules of knowledge. To deliver a TED Talk, however, is the apex of self-branding. And, not coincidentally, one of the major ways people discover TED Talks and other self-promotional videos is through Facebook.[8]

On Facebook, everything is an advertisement and advertisements are everything. There are no clear markers that distinguish the lines between items meant to persuade (to vote, to buy, to donate, to play) and items meant to entertain or inform. Even the style of presentation of individual profiles function as advertisements. Facebook scrambles the commercial and the social. Our interaction with Facebook reinforces this. We are frequently reminded that we must groom our social media profiles to offer the most compelling and pleasing image to the world. If we say the wrong thing or post the wrong photo on Facebook, we could find ourselves out of a job or a relationship. The scrambling of our social contexts creates constant pressure to perform and perfect our Facebook selves. We are constantly trying to persuade others on Facebook. We groom our profiles like we groom our hair. We choose what we post to perform our identities and affiliations. It's easy and possible to declare to the people you most want to

impress or commune with, "I am the sort of person who cares about Real Madrid, free-market economics, and the European Song Contest." Facebook profiles are advertisements for ourselves. Facebook sorts users' profiles by affiliation and interest, and then more dynamically based on interactions and the flow of "engagement." And users feed clues into Facebook's algorithms that amplify that process further, creating a constantly churning lattice of affiliations.[9]

The more time we spend grooming our self-presentation and declaring our tribal affiliations, the more we grow acculturated to this habit. It becomes a cultural norm. It assumes the soft power of ideology. Companies that desire to break through the flows and distractions all around us find it necessary to detect and play to those affiliations. If I have repeatedly declared myself a fan of the San Antonio Spurs, companies that wish to sell me things would like to know that. To grab my attention away from all the things my Facebook Friends post about the Spurs, both for and against, advertisers have to hook me with something fresh and loud so that they can move me to consider buying something related to my interests. This is not only possible with the massive collections of data about each of us that marketing firms, Facebook, Google, Amazon, and others possess, but it is efficient and inexpensive.

We may have reached "peak advertising." If our information ecosystem becomes completely saturated with attempts to pull our attention toward products and services, we might build up resistance to each effort. News, entertainment, web search, and even social networks would find their marginal returns for each dollar invested in advertising slipping away. If our attention becomes so completely harvested that there is no more for us to give, how would we have dependable and affordable news? There would be no space in our world unmarked by commercial sponsorship. Traditional advertising has had a great run for almost two hundred years. While its practices and technologies have shifted significantly over that time, the basic functions remains the same: grab attention from potential consumers, focus it on a good or service, and have the provider of that good or service pay for that transaction to subsidize whatever media content the consumer had hoped to encounter. The consumer does not have to pay much up front

for the content. The media company gets diverse sources of revenue. And the firm that buys the ad gets some faith that it can locate potential consumers efficiently. That's how it's supposed to work. But as each advertisement loses its power in the torrent of draws on our attention, more ads flood our field of vision.[10]

Many public school systems across the United States now allow corporate advertising to target students in exchange for a flow of shared revenue for the schools. Lockers, hallway walls, and meals bear logos. School announcements appear in video billboards amid advertisements. In America in 2018 few spaces are considered sacred or safe from the demands for our attention. The advertisements that seem to work best are those placed through services that precisely target the good or service to the person who might want that good or service, and those that seems least like an advertisement. Facebook accomplished both of these by erasing the cues that help us distinguish between a commercial for a product and a post from one of our Friends. And it accomplished that most powerfully by harvesting data about our behavior and preferences to ensure we see ads that Facebook's algorithms judge to be "relevant" to us.[11]

This is a radical change in how advertising works. Data-driven advertising upends almost two centuries of advertising that's based not on empirical evidence but on a certain faith that the message will reach its intended audience; this earlier type of advertising funded the media systems that enabled many of the changes of the twentieth century. The mania to capture our attention began with the "penny press," which radically altered journalism and commerce in North America. When some nineteenth-century newspaper publishers realized they could sell papers for less than the *New York Times* charged per issue and make up the difference in bulk sales, they traded subscription and newsstand revenue for advertising revenue. All they had to do was demonstrate to local stores that they had a stronger grip on the attention of the rising literate middle class. This move, which never destroyed the *New York Times* yet was led by papers long ago preserved in the amber of library microfilm, served as the model for similar insurgent attacks in other media through the next 180 years. Despite its failure to destroy its rivals, the penny press movement did alter the practices and

expectations of all newspapers in America and—more significant—served as a model for how other media forms would deploy advertising in their efforts to reach broad audiences with minimal up-front charges to consumers. Desires to capture and measure consumer attention conflict with dreams of providing the sort of substantive, high-quality culture and information that a thriving democratic republic and dynamic economy need. But overall, advertising has managed to keep alive publications as diverse in depth and seriousness as the *Daily Mirror*, *National Enquirer*, *Economist*, and *New Yorker*. Reading markets, like consumer markets, are segmented. Advertisers and publishers have become masters at targeting both content and ads to particular tranches of the American public.[12]

While newspapers and magazines—and thus their advertisers—had long focused on broad segments of a mass market, generally marked by location, ethnicity, gender, and social class, by the 1970s the rise of computer databases and more subtle market research led to some new market segmentation models. These new segments were defined by cultural style as much as gender or ethnicity. Segment names like "Bohemian Mix," "Shotguns and Pickups," and "Young Suburbia" now helped advertisers, marketers, and the publications and programs that depended on advertisers reach audiences with new precision. This data-intensive (for its time) work signaled the end of the assumption that there is a single mass market for just about anything in America. It began the process of the "sorting" of each of us into pockets of like-minded market segments.[13]

In a steady climb over the next forty years, media, marketing, and advertising firms tapped larger pools of data, using more powerful algorithms and processors. Through direct mail they were able to target smaller and smaller tranches of consumers. But the loudest and most persuasive voices were still major publications, commercial radio broadcasts, and television. Market segmentation on television could really only get as precise as a region. So one advertisement for a product might run in Salt Lake City, Utah, while another version for the same product might run in San Francisco. Advertisers rarely employed even that level of segmentation on television because producing and testing targeted ads was expensive and time-consuming. It was more cost-effective to play to the presumed large middle of an audience.[14]

Internet-based advertising offered the ability to label and sort audiences with precision beyond even direct mail. Now advertisers could reach specific members of a household without mailing a piece of paper to the entire household. More important, web advertising creates data that show if the ad actually worked. All other forms of advertising, no matter how scientifically targeted or tested, are leaps of faith. But the advertising system Google invented in the first few years of the twenty-first century operated completely differently. Advertisers would pay for an ad only if a person clicked on it. This meant that if an ad failed, the advertiser could try a different version or target it differently. Digital media allowed that malleability. It's easy to experiment with different ad content, different audiences, and different targeting keywords. Also, Google assigned prices for ads based on instant, computer-driven auctions. The high bidder would pay only the price that the second-highest bidder had offered, ensuring that the auction winner never overpaid for an ad. After Sheryl Sandberg left Google to run Facebook in 2008, Facebook adopted many of the same techniques that had made Google valuable. Then Sandberg directed new experiments and ad services that only Facebook could provide, given the richness of personal data and the graph of social connections. This is the main reason Google and Facebook are the most effective advertising systems ever created. The 1960s and 1970s masterminds behind rudimentary market segmentation, inspired by the ethnic and cultural recognition movements of the 1960s, started a revolution that the Mark Zuckerbergs of the world have seen to fruition.[15]

Advertising in the mid-twentieth century focused on methods of persuasion. Experts in the field promised they could move or inspire people to purchase goods they might never have considered. They promised they could persuade undecided voters to move to support a candidate they might have ignored before. Steadily, as the twentieth century gave way to the twenty-first, advertising became much more about aligning interest with offers. One did not have to try to persuade consumers to buy a good or service if you could find enough new consumers who were already interested or—even better—had already searched for a product or service. Once an advertiser identified that consumer, it required less expense and

effort to urge the consumer to make the purchase. Direct marketing had long held out the dream of precise profiling to match the interest to the offer. Not until firms had deep records of searching, shopping, and personal interest behavior could they realize this dream. In an advertising world dominated by Facebook and Google, there is a lot less attempted persuasion. The game is all about matching.

The unrelenting drive to surveil and tag complex consumers and then demand slivers of their attention in hopes that they engage in a series of transactions is more than distracting and exhausting. It's dehumanizing. It treats us each as means to a sale rather than as ends in ourselves. It also undermines centuries of concerted effort to engage in complex collective thought that could enable us to deliberate about the issues we face as a species. It's hard to participate in a republic, let alone face global challenges, when hit network programs such as *The Voice* have our eyes darting from television to iPad to phone, tweeting and cheering and chatting and shopping along.

Not that there is anything wrong with that. Tweeting, cheering, chatting, and shopping are fun. And no one said we signed up for a republic without fun. The difference now is the unrelenting ubiquity of these draws on our attention. Even our political lives are multiscreen affairs. The professional heirs to radio visionary David Sarnoff's NBC empire have spawned the specter that now haunts the party of Lincoln, the constitution of Madison, and the nation of Jefferson—a reality-TV-star-in-chief.

HOW FACEBOOK TURNS PERSONAL DATA INTO PROPAGANDA

In 2016 agents working on behalf of the Russian government used the very Facebook feature that has made the company stunningly profitable and convinced so many businesses to shift their advertising budgets from newspapers to Facebook. The Russian agents targeted content using the Facebook advertising system to mess with American democracy. They created bogus Facebook groups and pages devoted to such issues as opposing gun control, opposing immigration, and pushing for Texas to secede from

the United States. Russian agents even ran one Facebook page called "Blacktivist," purporting to support a campaign against police violence. Once U.S. officials pressured Facebook to come clean about the extent to which the company had been hijacked by Russian operatives, Facebook found 470 pages and profiles linked to a Russian company called the Internet Research Agency. The people who controlled pages had purchased about three thousand ads, often paying in Russian currency. Ultimately these techniques reached more than 126 million Americans. These Facebook pages and groups managed to motivate more than 62,000 Americans to pledge to attend 129 rallies and events meant to support Donald Trump, oppose Hillary Clinton, and protest mosques around the United States. Each of these Americans saw these posts because Facebook helped the Internet Research Agency target them through the company's basic advertising service.[16]

Setting up an ad campaign on Facebook is both easy and inexpensive. It's self-service. And there is alarmingly little oversight from Facebook. Such an expansive global system can't easily be controlled. So Facebook leaves it largely uncontrolled. Its chief form of governance is machine learning. As advertisers generate categories of Facebook users to target, or Facebook users cluster through interaction, Facebook's algorithms create new categories to focus the reach of an ad. So there are many categories of targeted users from which to choose. To promote a podcast that I produce, I ran a campaign that spent only $200. I chose to push the notice for the podcast, which is about how scientists who study the brain and behavior do their work and discuss controversies, to a select audience in the United States, Canada, and the United Kingdom. I chose to focus the campaign on those who had expressed interest in psychology and neuroscience. I limited the ad placement to those who had an M.D. or a Ph.D. And I excluded those who were younger than thirty years old. This meant I would only reach about three thousand Facebook users. But they would be the right three thousand Facebook users. Just for fun I also excluded any Facebook user who had expressed interest in the 1970s country music singer Crystal Gayle. This did not change the size of my target audience at all. Or if it did, it reduced it by only a handful of Ph.D.'s who happen to like Crystal Gayle.

This ability to exclude types of people from an advertising campaign is crucial to the success of Facebook ads and therefore of Facebook itself. Only Facebook knows who in the United States, Canada, and the United Kingdom has expressed interest in Gayle. This might seem silly and harmless. It is. Facebook also, notoriously, has allowed advertisers to exclude Jews, African Americans, women, and Spanish-speakers from advertising campaigns, including those for housing and employment.[17]

Facebook's advertising services permit those running an advertising campaign to carry out tests on sections of the audience with different versions of the ad. If one version of the ad, perhaps with a red background instead of a green background, generates better "engagement" (more clicks, shares, comments, and likes), then the advertiser knows to drop the green version. Then the advertiser can test red against blue, or red against yellow, or red in California and yellow in France. Such constant iteration can make advertising campaigns better over time. Without adequate personal data on which to build such a targeting system, Facebook could not offer this service. Without such a targeting service, Facebook would just be running clumsy, ineffective ads—the same kind that newspapers and television stations have run for decades. This is yet another way that Facebook segregates us and creates alternative cultural and political universes for us. It's the core business of Facebook. To reform it is futile. It's Facebook.

FUNNEL VISION

The same logic favors some ads over other ads and some user-generated content (the photos, posts, videos, and questions we all post) over other content on Facebook. Engagement wins. Over time Facebook not only rewards the items that are likely to generate the most markers of engagement (clicks, likes, comments, shares) but also learns to tailor the News Feed of every Facebook profile. The longer you are on Facebook, the more you engage with items on Facebook, and the more you teach Facebook to send you more of the stuff that is very much like what you have already indicated interests you. Both Facebook and Google call this a test of "relevance." Relevance has no relationship to the helpful, the enlightening, the moral,

the educational, or the true. That logic works very well for product and service advertising. But when the same logic applies to all posts, whether they are news or news-like items or videos of cats on skateboards, the experience over time is one of a narrowing field of vision. We start to see more items posted by Friends with whom we have richly engaged before. We see more items posted from publications or entertainment outlets with which we have engaged before. It's more than tunnel vision. It's funnel vision.

This phenomenon is what writer, activist, and entrepreneur Eli Pariser dubbed the "filter bubble." Facebook gives us more of the sorts of items we respond to with clicks, likes, shares, and comments, pushing aside things that might not interest us. That's fine if the subject is commercial or trivial. Cat people should not have to see too many items about dog products and vice versa. Pariser noticed in the early days of Facebook that over time he saw fewer posts from Friends who differed with him politically. He suspected that Facebook was algorithmically clustering people, generating this funnel vision. So he interviewed engineers at Facebook and Google and listened to their explanations of the values they embedded in their algorithms. Pariser's suspicions were valid. Once his book *The Filter Bubble* (and the accompanying TED Talk) debuted in 2011 others began to doubt the impact of the phenomenon. The power of filter bubbles remains unmeasured and perhaps unmeasurable, despite many efforts to do so. The intent is clear, however. Both Facebook and Google declare that their algorithms perform the function of delivering "relevant" content based on our record of engagement and what others like us and around us do. The phenomenon is observable on a personal level, as Pariser and others have testified. But the question of whether the filter bubble actually limits our vision or tribalizes us even more than we already do for ourselves remains murky. It's almost impossible to measure the effects of the filter bubble without access to user data that only Facebook has.[18]

In 2015 researchers at Facebook published in *Science* a significant study of political polarization among Facebook users. The study, which can't be replicated or deeply examined because Facebook keeps its data to itself, concluded that "compared with algorithmic ranking, individuals' choices played a stronger role in limiting exposure to cross-cutting content." This

means that Facebook researchers found a significant amount of polarization among Facebook users but ascribe it to homophily—a sociological term used to describe our urge to cavort with those similar to ourselves—more than to Facebook's algorithms. It's easy to make too much of these findings. And many did. First, the paper does not deflate the idea that the algorithms sort items along patterns that reflect previous engagement. It just concludes that the algorithmic effect is lighter than the effects of the choices we users make.

The more important lesson to take from this study is that Facebook and user homophily work synergistically to narrow our field of vision. It's wrong to conclude that we have forged "echo chambers" that seal out divergent views, or that Facebook has. That's an extreme and unfounded assumption. But it's reasonable to believe that our vision is narrower than it might otherwise be if we were not engaging with Facebook so often and about such important matters. The choices we make and the choices Facebook makes for us feed each other. We are part of the system. The technology is not distinct from the culture, the politics, and the ideologies in which it operates. The arguments about whether the filter bubble exists or how much it matters too often get bogged down in a false dichotomy: is it the fault of technology or is it the fault of humans? The answer is always yes.[19]

Political segregation, the fraying of civil discourse, and the erosion of trust in civic institutions is a fact of life around the world. So the narrowing of our field of vision, the filtering of information to that which emboldens our confidence in the beliefs we already hold, and the segregation of citizens into polarized camps all steadily undermine the potential for republican governance and reasonable discourse about important matters.[20]

One of the reasons it's easy to doubt the power of the filter bubble is that some have been crying wolf about it since before the problem even emerged. The earliest critique of what became known later as the filter bubble predated Facebook by three years. In 2001 legal scholar Cass Sunstein reacted to a proposition then in vogue among internet idealists: once all of our news content existed on the World Wide Web, we could use various filters and platforms to curate a flow of news that we could customize to our interests. For more than a century we bought a whole newspaper just for

the sports section, the comics, or the bridge column. We got the news from Moscow and Melbourne in the package, whether we cared or not. Newspapers had no way of knowing if people read any of their stories from beginning to end. They had no way of knowing if anyone ever opened up the business section. Like the companies that advertised within them, newspapers operated on only the roughest data, total sales, and chose content based on hunches, trends, or even something called news judgment. The web offered readers a chance to encounter only content that they selected. So they might read four stories about the World Cup tournament matches and none about the corruption within FIFA. They could encounter many pieces about the struggles in Puerto Rico following a hurricane and nothing about the suffering in Mumbai after a massive monsoon. Big thinkers such as MIT professor Nicholas Negroponte celebrated this sort of customization. Sunstein was appalled. This reeked of consumerism trumping republicanism, as he wrote in the book *Republic.com*.[21]

Sunstein wrote the same argument with a different set of conditions two more times. In 2007 he updated *Republic.com* with *Republic.com 2.0*. The thing that had changed was the rise of blogs between 2002 and 2007. Then, in early 2017, Sunstein released his third version of the same argument, *#Republic*. This time, the focus was on how social media generated the "daily me," but the concern was the same. The central argument of all three versions is that our growing habit of relying on websites, blogs, and now social media has focused our vision and sorted us into what he calls "echo chambers," ensuring that we fail to see information that complicates our sense of reality and fail to engage with views that challenge our position. Sunstein deserves some credit for identifying this problem many years before other media critics realized it. But he lacked strong evidence for his concern until the rise of powerful algorithms such as those that drive Google and Facebook. The 2001 version of his argument challenged the rhetoric of web enthusiasts who proclaimed that the consumerist vision of a tailored media environment would be wonderful for both individuals and consumers. But in 2001 web pages were still static and fairly new, and the filtering power of Google was only just becoming clear. By 2007 Sunstein was compelled to focus his argument on the blogosphere, just about five

years old at the moment of publication. Blogs brought a social dynamic to the web's news, information, and commentary ecosystem through the then-influential blogrolls, which helped readers discover similar blogs recommended by their favorite bloggers.[22]

The problem for Sunstein's thesis was that many of the most influential political blogs at the time regularly and mostly respectfully engaged with ideological opponents in debates that were of high quality. Conservatives such as Glenn Reynolds, Ann Althouse, Mickey Kaus, and Andrew Sullivan mixed it up with liberals including Heather Digby Parton, Jack Balkin, Matthew Yglesias, and Duncan Black. A post on a liberal blog would challenge one from a conservative blog with a link and full quotes. The response on the conservative blog would link back to the liberal one. And each writer would use whatever space it took to outline an argument and offer evidence—also supported by links to sources. Just as important, the high moment of blogs allowed academic experts to exert public voices in ways that had been rare or impossible just a few years earlier. As sociologist Eszter Hargittai found in 2008, the most notable political bloggers across the political spectrum often addressed each other's writing substantively, whether to agree or to disagree.[23]

Two of those newly empowered academic bloggers, Daniel Drezner and Henry Farrell, responded to Sunstein's dour account of blogging with a full-throated defense in *Foreign Policy* of how valuable and important the blogosphere had been to enriching public understanding and discourse. Nonetheless, Farrell coauthored a paper in 2009 that concluded that blog readers tend to read blogs that accord with their political beliefs. Cross-cutting readership of blogs on both the left and right of the spectrum was rare. Farrell and others found strong evidence of polarization among blog readers, who tend to be more polarized than both those who do not read blogs and those who consume television news. Both of these things could be—and were—true: the major political blogs encouraged readers to engage with differing points of view, and regular readers of blogs sorted themselves into clusters by starting their explorations among a core set of blogs with which they agreed strongly. In the age of the blogosphere, no algorithm amplified homophily. But homophily still mattered. Readers

segregated themselves, which should be no surprise. We all love to remind ourselves how smart and correct we are. And we are all susceptible to confirmation bias, approving of accounts that reinforce our assumptions and rejecting evidence that runs contrary.[24]

We should look back on the high moment of blogging with nostalgic reverence. The blogging era, roughly 2002 to 2007, was perhaps the richest and most diverse media ecosystem we have ever enjoyed. Google was gaining ground in the advertising market, but the market had not yet moved so destructively away from freestanding publications, both web-based and print-based. Web-based sites such as OpenDemocracy.net and Salon.com grew in status and influence. Diverse opinions and analyses generated by informed and articulate experts found audiences through blogs and the links that appeared on blogs. One interesting blog would lead a reader to another.

Since then, though, Facebook and Twitter have drained the blogosphere of its vitality. Facebook and Google have drained publications of revenue and confidence. Some group blogs such as Huffington Post, Talking Points Memo, and Boing Boing have morphed into influential professional news and commentary sites. But the communal DIY spirit is gone. Much expression that used to fill blog posts now finds itself distilled in 140 characters on Twitter or posted on Facebook and thus subject to the whims of Facebook's powerful algorithms.

The Facebook algorithms made the difference. We chose our Facebook Friends. And most of us have friends, relatives, and colleagues who have different opinions and read different sources than we do. But our habits of interaction—clicks, comments, and "likes"—tell Facebook to reward the sources that earn our "engagement." We get more of the same. We are already creatures of habit and comfort when it comes to our views of the world. Facebook radically amplifies those tendencies. Things posted on Facebook can't easily travel beyond Facebook. And Facebook posts find only the audience that Facebook decides it deserves. Such "choice architecture" is dangerously powerful. We can believe that filter bubbles exist because Facebook and Google tell us what their companies do. They tell us—roughly and without precision or clarity—why we see certain content and

not other content. Basically, we tend toward homophily. Our media system can correct for that tendency or it can reinforce it. Our current media ecosystem, dominated by Facebook, amplifies both our weakness for confirmation and our desire to cluster among the like-minded. It's not just the technology. It's not just the cultural proclivities of human beings. Technology is culture and culture is technology. All are part of the same system.

FEEDING THE BEAST

Before Facebook rose to its central role in how people learned about the world around them and performed politics, Google's search algorithms largely determined what got read or seen. Sites that optimized their design and content to satisfy Google's search algorithms could thrive. When people turned to Google as news broke about a mass shooting or a tsunami, the sites that could please Google would earn the clicks. The clicks would justify advertisements. And more clicks and links from other sites would mean more visibility in the future. Success in the "search engine optimization" game could generate positive feedback loops—uphill cascades—so the winners would keep on winning and those sites unable to pander to Google would fall to the dark depths of the second page of search results.[25]

Out of the flurry of homemade and bespoke blogs that proliferated in the first decade of the twenty-first century, Huffington Post, launched in 2005, was the behemoth. Arianna Huffington seemed to be the ideal person to invent a new style of publication for the attention economy. She had, after all, managed to generate attention across multiple media forms (books, magazines, television) about multiple issues (modern art, conservative politics, liberal politics, lifestyle fads). She had written a controversial biography of Pablo Picasso. She had married a millionaire who lost a race for governor of California. She had appeared on talk shows of all sorts. But, having exhibited no significant mastery of anything, no stable ideology, and no business acumen, Huffington was most famous for being famous. Investors gladly bet that the attention economy would reward the sort of person who knew how to make people look and listen. They were not

wrong. Huffington's idea was to set up a sort of group blog. Her celebrity friends would contribute posts. People would flock to read what Ted Danson had to say about the state of our oceans or what Larry David had written about the U.S. ambassador to the United Nations. Just as important, the publication that would be dubbed Huffington Post (now Huff Post) offered unknown or yet-to-be-known writers a chance to contribute articles for no remuneration. By attracting writers who did not expect to be paid, Huffington kept expenses low. To keep interest high, Huffington depended on two partners who had already demonstrated a mastery of the new digital environment and who would soon make their mark in the next era of digital media: Andrew Breitbart and Jonah Peretti.[26]

Breitbart was a conservative culture warrior. He had worked for right-wing gadfly Matt Drudge and had mastered the style and substance that made the Drudge Report required reading for the politically obsessed beyond even its core conservative audience. Breitbart, who died in 2012, understood how to push emotional buttons and rile up members of a movement. He had a trickster streak. And he was frenetic and obsessive, making him an uncomfortable business partner for the Huffington Post venture. Soon after the launch Breitbart left Huffington Post to launch his own site, one that would embody his politics and personality: Breitbart.com. That site, taken over after Breitbart's death by another frenetic right-wing populist, Steve Bannon, would master the social media era of the attention economy much the way Huffington Post had succeeded in the search engine optimization period.[27]

Peretti stayed at Huffington Post until AOL purchased it in 2011, but in 2006 he launched what he called a "laboratory" to test how content might flow as the media ecosystem changed. He called that laboratory BuzzFeed.com. Over the past decade BuzzFeed has grown from a source of viral lists, quizzes, images, videos, and cats to a source of viral lists, quizzes, images, videos, cats, and serious, lively journalism. BuzzFeed has investigative journalists. It has bureaus around the world. It has top-notch cultural commentators. And it still has listicles and cat photos. What BuzzFeed learns about testing, revising, measuring, and promoting posts of one kind helps it distribute posts of other kinds. It's a giant learning machine as well

as a news and entertainment machine. If any news organization has cracked the code of Facebook, it's BuzzFeed. And by specializing in ethically suspect paid content (advertisements that work like and look like editorial content) BuzzFeed has made a bet on the sort of advertising that Facebook and Google can't and won't do, specially tailored to BuzzFeed readers and distributed via the same algorithmic tricks that Facebook plays on the editorial content. It seems to be working, although the massive tide of advertising moving to Facebook and Google and away from other sites and news organizations has also caught BuzzFeed in a financial bind. If BuzzFeed can't make it in the long run in an information ecosystem dominated by Facebook, nothing can.[28]

Both BuzzFeed and Breitbart are at once tentative success stories and cautionary tales. Breitbart is the source of virulent ethnic nationalistic propaganda that aims to undermine the civic spirit of the United States. It's influential with many people who support President Donald Trump and who work around him. Material that spews out of Breitbart works its way up to cable news shows, often to be denied, debunked, or reported on as a "controversy." So Breitbart performs a disruptive and distracting role in the public consciousness. What Breitbart produces plays very well on Facebook. The people who run Breitbart learned the same lessons that informed Peretti and his associates at BuzzFeed. Separated at birth, both gain and maintain an audience by working with and through the ways Facebook energizes some posts and stifles others.

Their stories also should make us cautious because they signal to other news and entertainment outlets that one should—or must—pander to Facebook to succeed in a shrinking market for advertising revenue and a crowded supply of attention-seeking destinations. Not only do the *Guardian*, *El País*, and *Haaretz* all compete with HuffPost, Breitbart, and the *New York Times* for space and frequency on Facebook News Feeds, they also compete with YouTube videos, games, music, podcasts, and hundreds of other diversions in daily life—all of them more and more precisely engineered to hook us and keep us coming back.

By the summer of 2017, the Pew Research Center revealed, 67 percent of adults in the United States used social media to discover news. That was up

from 49 percent just five years earlier. While this is a rough number and might not reveal that much, largely because it's unclear what American social media users consider "news," the upward trend is strong and thus worth considering. Clearly social media and the "choice architecture" that governs them are becoming more important in our lives as citizens as traditional sources of news slowly go broke. To sum up, news organizations are compelled to make editorial choices (which stories to cover, how to craft headlines, whether to favor video over photos or text) to pander to Instagram's or Facebook's algorithms because they can't afford to be invisible on Facebook or Instagram. Yet advertisers flock to Facebook and Instagram and away from the sources of all that essential information about the world. Journalism is feeding the beast that starves it. And there is nothing journalists can do about it.[29]

THE OPERATING SYSTEM OF OUR LIVES

Soon we might not have to pay attention at all. If the five biggest technology companies in the world fulfill their wishes, they will sell us a series of devices, each promising to make our daily tasks a little bit more convenient. These devices would monitor us without our direct interaction. They would capture our intentions and desires at those few moments when we are not staring at screens and typing on keyboards. Some of these devices sit on our counters. Others are embedded in our cars. Some are built into our thermostats and appliances. Others sit on our skin. In 2016 the most successful advertising company in the history of the world introduced an obelisk called Google Home. It's a loudspeaker and a microphone constantly connected to Google's servers and algorithms. Always on, it passively listens to the words and sounds of the home. Google Home melds the information it gathers with the rich dossier that Google already has on its users. It promises to provide personalized responses to verbal search queries and requests to play music. Google Home is a response to Amazon's Echo, which also sits on a counter, constantly listening for voice commands to order more products for quick delivery. The newest version of Apple's Siri is also capable of listening at all times, learning constantly about its owner's

desires. Between these devices on our counters and the mobile sensors we increasingly strap to our skin and carry in our pockets, the "attention brokers" need no longer compete for our attention.

Consider the five most powerful digital companies in the world: Facebook, Alphabet (the holding company for Google), Microsoft, Amazon, and Apple. Through most of 2017 these companies were the five most highly capitalized companies in the world. We call them leaders of the "technology sector." Yet they are very different companies that do different things for us. While they compete for highly trained labor, within commercial markets they only compete incidentally. Facebook and Google do compete for web and mobile advertising. But they harvest and sell our attention in different ways because we use these services in different ways. Apple thrives by selling hardware for its core business. Microsoft still sells and leases software to businesses and consumers to make most of its money. Amazon is a retailer of goods that makes most of its money providing web server services.[30]

Yet these companies do compete in a larger, more visionary sense. Each of them has mounted efforts to achieve dominance in the next frontier of digital commerce: the data streams that would monitor, monetize, and govern our automobiles, homes, appliances, and bodies.

They no longer strive to be the operating system of our laptops and desktops. That competition is static and divided between Microsoft and Apple. They don't wrestle as much to be the default operating system of our mobile devices. The duopoly of Apple's iOS and Google's Android has locked out all other competitors, and both firms are thriving in that market. All five of these companies share one long-term vision: to be the operating system of our lives.

The operating system of our lives would measure our activities and states of being and constantly guide our decisions. Data, these visionary companies predict, will soon flow from our clothing, our vehicles, and thus our bodies. Most often this vision of data flowing through everything, with everything reporting back to central, controlling, algorithmically governed computers, is called the "internet of things." But that's not the right term. This new network is about people, not things.[31]

As they steadily introduce personal assistants, new interfaces, self-adjusting thermostats, self-driving cars, internet-connected glasses and watches, and virtual reality goggles, these companies hope to earn the trust of consumers and regulators so that they can set the standards for the transactions that can make this operating system work seamlessly and efficiently. If they realize this vision, there would be no clear distinction between media and non-media. There would be no distinction between content and objects. All objects and all bodies would be mediated content.

Such a prospect should alarm us. How shall we confront such an effort? How shall we muster a vocabulary, a set of theories, and a set of responses sufficient for such a challenge? For one, we should move away from discussing the "internet of things." The abstract vision of "the internet" is that of a network of networks, a set of connections that together create an open, accessible, dependable platform on which others may freely build and send messages without restrictions. The fact is, "the internet" has been a myth all along. The internet has never been a global, open, distributed "network of networks" that can connect all of humanity. In most of the world, digitized network communication is not so open, not well distributed, and not necessarily run through the sort of computer networks that serve as the foundation of internet communication in the United States. When you send a message via the AT&T mobile network, it's not really going through the internet of our dreams. It's going through a highly regulated proprietary system. If you are using an iPhone, that message is going through a highly regulated proprietary device. Because we increasingly send data from and to these machines that sit on our bodies and travel with us throughout the day, these flows of information are not part of some separate sphere or "cyberspace." The basic architecture of the operating system of our lives would be even more closed and restrictive. Firms and governments are quickly learning the hazards of connecting thermostats, robots, cars, and watches to a fairly open network. These networks are insecure and the humans embedded in them are undependable. Malevolent political movements, criminals, hostile states, or pranksters can hijack or bring down a delicate system of data flows. So as the architecture of this system moves on, security will be a priority over interoperability. Data networks are likely

to be proprietary and access to them tightly controlled. For such a system of connected objects and bodies to operate properly it would have to be controlled by one company or a conglomeration of a handful of companies. This prospect of concentrated power should overshadow current concerns about the duopoly of Facebook and Google. Whichever company comes to dominate the operating system of our lives will wield massive power and will be an opaque "black box," governing many aspects of life with no transparency or accountability.[32]

We misunderstand this tangle of digital technologies and the roles they play in our lives when we assume that "the internet" has a particular logic, must necessarily have particular attributes (openness, neutrality, etc.), or constitutes a special "place" that operates by its own rules and norms, rather than amplifying some of the rules and norms that we already hold and undermining others. More significant is that we misunderstand the effects of technologies when we pretend that they operate or exist outside of human bodies and human relations. We are all embedded in the data network that would constitute the operating system of our lives. And the data would flow through us. Like with any technological system, we shape these technologies as much as—if not more than—they shape us. And by thinking about "the internet" or the "internet of things" instead of specific technologies and companies and how they interact with our minds and bodies, we miss the real trees for the imaginary forest.

Of these five companies, Facebook is the most pervasive and most dangerous. Facebook plays a peculiar role in the race to become the operating system of our lives. It is the most influential media company in the world. It shapes the messages that politicians, dictators, companies, religions, and more than two billion people wish to send to the world. It increasingly serves us news content, or content that purports to be news. It is the most powerful and successful advertising system in the history of the world. It's increasingly the medium of choice for political propaganda. While Facebook has thus far failed to develop devices that have invaded our homes or attached themselves to our bodies, it has assumed the goal of managing our personal lives, our commercial lives, and our political lives from within its constellation of services and applications. Across the world,

more people spend more time during their days staring at their mobile phones, interacting with Facebook, Facebook Messenger, Instagram, and WhatsApp. The data people generate by interacting with these Facebook-owned services shape them. And the design choices that Facebook engineers make for these applications shape how we see and deal with the world. That complex set of feedback loops, the rich flow of personal data, and the way that the ubiquity and utility of Facebook services feel almost given have put Facebook well ahead of the other four companies in the race to be our operating system.

Of the seven most-used social media platforms, Facebook owns four—Facebook, Messenger, WhatsApp, and Instagram. If we consider only social media services that operate outside China, Facebook owns four of the top five; only YouTube, owned by Google, occupies the non-Facebook spot, at number two. Twitter, often erroneously considered a peer to or in competition with Facebook, ranks a distant tenth among all the social media platforms. Twitter has only 350 million registered accounts (compared to 2.2 billion for Facebook in 2018), but some significant portion of those accounts are "bots," or automated accounts that amplify propaganda. In terms of actual influence on actual human beings, there is no digital or social media company with the reach and power of Facebook. It is unique in human history. Yet Facebook has a direct rival. Like all true rivals, it is also an inspiration. Facebook is not the most successful company at grabbing, holding, and processing human attention. That prize goes to Tencent, a China-based firm that makes the most successful social media application in the largest market on earth: WeChat.

With almost a billion users, WeChat has infused itself into their lives in ways Facebook wishes it could. WeChat offers photo and messaging services that resemble Facebook, Instagram, and Twitter. WeChat also performs web search functions similar to what Google offers. Beyond information and communication, WeChat users can use the service to purchase items from vending machines, pay utility bills, exchange gifts, scan QR codes, conduct bank transactions, schedule doctor's appointments, search for and reserve library books, and play augmented reality games such as Pokémon Go. A WeChat user can perform many essential daily tasks without looking

up from a phone or even switching among applications. WeChat is the model of one application becoming the operating system of our lives.[33]

Facebook has been building WeChat-like services into its various applications. And Mark Zuckerberg openly pines for an opportunity to operate once again in China so that Facebook can compete directly with or meld services with WeChat. When we consider the role that all those human decisions play in the WeChat or Facebook system, we must recognize that we are being mediated as much as any items of cultural expression or information.

Social media are no longer just social media. All media firms and services are attempting to become "social." Social bonds are "sticky," so they help in fending off others who would try to steal away attention. Social connections are also valuable for segmenting a market and targeting a message. As Facebook becomes more videocentric and superstores such as Amazon move into the business of producing entertainment, media forms converge. Observations and accounts of the political economy of Facebook in 2010 might not be relevant in 2020. What if "social media"—and in fact "media"—become something much more pervasive and powerful in our lives over the next twenty years?[34]

By attaching data-transmitting devices such as phones and watches to our hands we already have made ourselves nodes in the networks, subscribing to one or more companies' efforts to become the operating system of our lives. It could get more intrusive. There is a way for companies to record even more precise measures of our interests and preferences—even our fantasies. And it could grant companies greater feedback control over us in a field that has more addictive stimuli than even the best smartphone. In 2014 Facebook purchased Oculus Rift, one of the more developed virtual reality companies aimed at consumers. Zuckerberg has set a goal of getting one billion people to immerse themselves in virtual reality through Oculus. Toward that end, in 2017 Facebook released a new headset for the service that costs about $200, far less than any previous device. Two days before the company released the device, a cartoon version of Zuckerberg visited virtual Puerto Rico while its residents were suffering greatly after the floods and destruction from Hurricane Maria. The company broadcast

a live video stream that managed to combine two ethical violations: disaster tourism and product promotion. He appeared as an avatar that "teleported" to different locations using Facebook's "social VR" tool Spaces. All the while Zuckerberg was sitting at the company's Menlo Park, California, headquarters, dry and safe, with ample electricity flowing around him, wearing an Oculus headset. Zuckerberg half apologized for that stunning lack of compassion by claiming he was trying to "raise awareness" and was motivated by concern for the people of Puerto Rico. The sad truth is that while he *was* concerned, no one in his own ideological filter bubble raised any objections to such a stunt. Zuckerberg, as we will see in Chapter 4, makes no distinction between what is good for Facebook and what is good for humanity. To Facebook, everyone outside the company is a means to demonstrate the world-changing potential of the next big product the company offers. And Facebook, he believes, is a means to make everyone's life better. If he means well, Zuckerberg assumes, he must be doing good.[35]

"The idea is that virtual reality puts people first," Zuckerberg said when he first demonstrated Oculus Rift in 2016. "It's all about who you're with. Once you're in there, you can do anything you want together—travel to Mars, play games, fight with swords, watch movies or teleport home to see your family. You have an environment where you can experience anything." Once you're in there, Facebook would not only track every thought and gesture as you travel to Mars or fight with swords but also track all your social engagements through virtual reality, as it already does through the Facebook app. It would learn even more about you than it already knows. It would push and prompt and move you to feel certain ways and buy certain things more effectively than it already does. Virtual reality is more than a potentially mind-blowing technology for immersive simulation. It's the ultimate Skinner box. You might think you are playing a game with deeper involvement than ever before. But that's what every pawn experiences on every chessboard.[36]

Facebook already monitors our movements and motivations. Through virtual reality it hopes to capture our voices, fantasies, and body movements without our active compliance. If Facebook becomes the operating system of our lives, we could ignore it and it would still respond, monitor,

record, profile, sort, and deliver data—and more. We have a long way to go before humanity is constantly connected to and monitored by a handful of companies. Amazon Echo, Google Home, and Oculus Rift are currently just vanity products for the wealthiest among us. But the model is clear: the operating system of our lives would be about our bodies, our consciousness, our decisions. Attention would be optional. Power would be more concentrated, and manipulation constant. That's a world with no patience for autonomy and no space for democracy. It would be a lazy, narcotic world. It would not be some dystopian state of mass slavery, as portrayed in *The Matrix*. It would be kind of dull and kind of fun. It would be an existence like that in *Brave New World*.