

# The Legend of the Echo Chamber

IT IS 4:30 P.M. Dave Kelly has just finished his workday at an advertising firm in early September, 2018, and pops a CD into the stereo of his aging car. He is preparing to do battle with a formidable enemy: the New Jersey Turnpike at the beginning of a holiday weekend. When Dave finally reaches the exit for his hometown more than one hour later, he stops to perform a weekly ritual. Each Friday night, Dave checks out half a dozen books from his local library, cracks a can of overpriced craft beer, and settles in to read for at least an hour. This week he has chosen a mix of well-thumbed paperback novels, a book about the latest advances in cancer research, and a thick tome on human nature by an evolutionary anthropologist.<sup>1</sup>

Though he might not fit the stereotype of Donald Trump supporters, Dave voted for the former real estate magnate in 2016. Raised in a family of moderate Democrats, Dave veered toward the right in the 1980s because he was so impressed by the leadership of Ronald Reagan. But Dave is not a card-carrying member of the Republican Party. He cast two ballots for Bill Clinton in the 1990s, and takes liberal positions on most civil rights issues. “I’m perfectly happy with gay marriage,” Dave says. “I don’t understand why you would want to make an issue out of that.” But on economic matters, Dave is more libertarian. When he learned that New York City officials were considering a new law that would require businesses with more than five employees to

## CHAPTER 1

provide two weeks of paid vacation, Dave warned, “There’s gonna be a lot of companies that fire people to get away from that. There’s gonna be companies that just can’t do it and are gonna go out of business.”

Living outside liberal Philadelphia—and working in a profession dominated by Democrats—Dave normally hides his conservative views. “I have friends I won’t discuss this stuff with,” he says, “because I’m not going to change my mind and they’re not going to change theirs—so what’s the point?” The few times he tried to start such conversations, he explains, things quickly became heated—and the only thing Dave hates more than New Jersey traffic is heated arguments about politics. Because he feels like an unwelcome minority in his day-to-day life, Dave describes social media as a kind of refuge. He originally joined Facebook and Twitter to escape politics and follow updates about his favorite television shows. But he kept finding himself getting “sucked into political discussions.”

Over the past few years, Dave—who does not use his real name on social media—has spent many late nights arguing with Democrats on Twitter. Remembering one of these conflicts, Dave said, “Don’t judge me . . . I had a couple of beers.” A local radio station, he explained, had reported a group of White supremacists were planning to march on the campus of a nearby university. “Turns out they’re not,” he says. “The whole thing is a hoax.” After reading more about the story, Dave learned that one of the groups that raised the alarm was the progressive Southern Poverty Law Center. “They pretty much claim anyone who’s to the right of Karl Marx is a hate group,” he says. When he dismissed the incident on Twitter, another user quickly fired back, calling him a racist. “I called her an idiot,” he says. She didn’t know what she was talking about, he decided, because she was only getting one side of the story.

But so is Dave. Though he prides himself on being informed, Dave gets his news from a conservative talk radio station, the

## THE LEGEND OF THE ECHO CHAMBER

right-leaning website *Daily Caller*, and Twitter. Of the several hundred accounts that he follows on Twitter, only *New York Times* columnist Bret Stephens could be described as “centrist.” Dave has consumed a steady diet of conservative views on social media for years. Each day, his feed gets filled with content from Fox News, posts by Trump and other prominent Republicans, and dozens of memes bemoaning liberal hypocrisy. Dave has even retweeted a few messages from Russian trolls masquerading as American conservatives along the way. And that drunken Twitter argument about the White supremacist march at a local university? It turns out that Dave used more colorful language than “idiot” to describe his liberal opponent that night.

### The Echo Chamber about Echo Chambers

You might think you already know what’s going on here: Dave is stuck in an echo chamber.<sup>2</sup> Social media sites allow people to choose what types of information about politics they want to expose themselves to—or learn what Justin Bieber ate for dinner last night. The problem is that most people seek out information that reinforces their preexisting views. We connect with newspapers, pundits, or bloggers who share our worldview. If you’re a conservative like Dave, you might follow Tucker Carlson, the Fox News host, since you appreciate what he has to say about government spending or illegal immigration. And if you’re a progressive liberal, you might follow CNN’s Don Lemon because you appreciate his frequent posts about the issues you care about—racial inequality, perhaps, or climate change.<sup>3</sup>

The problem, the story goes, is that our ability to choose what we want to see traps us inside echo chambers that create a kind of myopia. The more we are exposed to information from our side, the more we think our system of beliefs is just, rational, and truthful. As we get pulled deeper into networks that include only like-minded people, we begin to lose perspective. We fail

## CHAPTER 1

to recognize that there are two sides to every story, or we begin listening to different stories altogether. Echo chambers have their most pernicious effect, common wisdom suggests, when people like Dave are unaware of them: when people think that they are doing research about an issue, but they are actually just listening to what they want to hear. When we encounter people from the other side, their views can therefore seem irrational, self-serving, or—perhaps most troubling—untrue. If we could only step outside our echo chambers, many people argue, political polarization would plummet.

The concept of the echo chamber existed long before social media did.<sup>4</sup> Political scientist V. O. Key introduced the concept in the 1960s to describe how repeated exposure to a single media source shapes how people vote.<sup>5</sup> The concept gained major traction, however, with the rise of 24/7 cable news stations in more recent decades. Social scientists quickly realized that such stations were allowing Democrats and Republicans to perceive starkly different versions of reality.<sup>6</sup> A popular example of the echo chamber effect is the 2002 U.S. invasion of Iraq. During this period, Fox News repeatedly claimed that Saddam Hussein, the Iraqi dictator, was collaborating with Al Qaeda, the terrorist organization responsible for the September 11 attacks. It was later discovered that such claims were false. But an influential study found that Fox News viewers were two times more likely to believe that such links existed than those who got their news from other sources.<sup>7</sup> If you are a Democrat, don't pat yourself on the back too quickly. A recent study showed more Democrats are trapped inside echo chambers than Republicans.<sup>8</sup>

Concerns about echo chambers gained added urgency with the rise of the internet and social media. In his influential 2001 book, *Republic.com*, legal scholar Cass Sunstein warned that partisan websites and blogs would allow people to avoid opposing views even more efficiently than cable news.<sup>9</sup> The internet activist Eli Pariser pushed this argument even further in his 2012

## THE LEGEND OF THE ECHO CHAMBER

book, *The Filter Bubble*.<sup>10</sup> He argued that algorithms employed by large technology companies made the echo chamber effect even worse. Facebook, Google, and other giant corporations exacerbate our built-in tendency to seek information that is aligned with our worldview via algorithms that recommend even more of such content to us. The most dangerous part of these algorithms, Pariser argued, is that social media users are not aware of them. Filter bubbles can preclude the very possibility of bipartisan interaction, Pariser warned, allowing our deeply biased views to go unchallenged.

Meanwhile, social scientists began to uncover substantial evidence of social media echo chambers as well. A 2015 study by data scientists at Facebook estimated only one-quarter of the content that Republicans post on Facebook is ever seen by Democrats, and vice versa.<sup>11</sup> A study of Twitter reached similar conclusions. More than three-quarters of the people who retweet—or share—a message, the study concluded, belong to the same party as the message’s author.<sup>12</sup> These findings were particularly concerning since social media was rapidly becoming one of the most popular ways for Americans to get their news. Between 2016 and 2018, the number of people who got their news from social media surpassed those who learn about current events from print newspapers. By 2018, social media had become the most popular news source for people ages 18–29.<sup>13</sup>

It should come as no surprise, then, that a growing chorus of technology leaders, pundits, and policy makers now warn of a grim future in which any discussion of politics on social media will quickly devolve into tribalism. We hear calls for social media platforms to break our echo chambers—or at least revise the algorithms that reinforce their walls. And if social media companies won’t relent, then social media users should begin stepping outside of their echo chambers themselves. Only then, many people believe, can we begin the difficult conversations needed to beat back polarization on our platforms.

## CHAPTER 1

It's a compelling story—especially when the people who tell it are those who helped build social media platforms and now regret their actions. But I believe the common wisdom about social media, echo chambers, and political polarization may not only be wrong, but also counterproductive.

### A New Lens on Polarization

Common wisdom often becomes unassailable because it is very difficult to verify.<sup>14</sup> Social scientists have wondered whether echo chambers shape our political beliefs for decades, but studying this process is very challenging.<sup>15</sup> We can analyze people like Dave Kelly—the craft-beer-drinking Trump voter described above—but are his experiences typical? Echo chambers result from the coordinated behavior of millions of people across sprawling social networks that evolve in complex patterns over time. Even if we had the time and resources to identify thousands of Dave Kellys—and see that people like him develop increasingly partisan views over time—how could we be sure that people's echo chambers shape their political beliefs, and not the other way around? If our political beliefs guide how we try to understand the world, would we really give them up so easily? Would Dave Kelly begin to moderate his views if we suddenly began exposing him to social media posts from progressive groups like the Southern Poverty Law Center?

Regardless of what you think about echo chambers, Facebook, Twitter, and other social media platforms have produced exciting new opportunities to study them. The social sciences were once considered “data poor” compared to other fields of study. But some platforms now allow us to collect information about millions of people in seconds. Even more importantly, we can now conduct an epidemiology of ideas, tracing how beliefs about the world spread across large social networks over time. The age of computational social science—the study of human behavior

## THE LEGEND OF THE ECHO CHAMBER

using large digital data sets—also provides new opportunities for experimentation. By embedding randomized controlled trials within social media platforms, computational social scientists have been able to increase voter turnout, organ donation, and a host of other positive human behaviors.<sup>16</sup> These types of experiments also hold enormous power to provide insights into social media echo chambers, as we will see.

But there is also a dark side to computational social science. In 2013, the psychologist Michal Kosinski launched a study to determine whether patterns in social media data—such as information about the things we “like,” or the accounts we follow—could be used to predict our ethnicity, sexual orientation, or even our intelligence.<sup>17</sup> Kosinski and his team produced an app that allowed Facebook users to perform a personality test on themselves via the data generated within their accounts. But the now-infamous political consulting firm Cambridge Analytica allegedly created a similar app to collect data for a nonacademic purpose: creating microtargeting campaigns to sway political elections.<sup>18</sup> Though many social scientists question whether such ads were effective, the story highlights a dangerous precedent: the tools of computational social science can be repurposed to violate privacy and potentially manipulate the behavior of people who did not consent to be studied.<sup>19</sup>

Computational social science has another problem too: the digital footprints we leave behind on social media platforms provide a very incomplete record of human behavior.<sup>20</sup> As a thought experiment, let’s put Dave Kelly’s data into the type of app created by Cambridge Analytica. We could easily conclude that Dave is a Republican by analyzing the news organizations and pundits he “likes” or “follows.” A political campaign might even be able to identify which television shows Dave watches and buy commercials to reach people like him. But we would also misunderstand some of the most important things about Dave. Though his Twitter feed makes him seem like an angry “Make America

## CHAPTER 1

Great Again” warrior, the app would not reveal that Dave is actually worried about climate change and disappointed by his party’s treatment of gay people. You’d never know that Dave thinks Trump is a bully, or worries about racial discrimination in policing. You would not learn that Dave was skeptical about whether White supremacists were really marching at a nearby university during the incident I described at the beginning of this book because he believes media organizations are stoking ethnic tensions for financial gain. Most important, you would not learn that this issue is particularly important to Dave because he is part Puerto Rican and suffered terrible discrimination as a child. I mention these details not only to show how many things are left out of the digital record of our lives. On the contrary, I believe the rapidly growing gap between social media and real life is one of the most powerful sources of political polarization in our era.

How did I come to this conclusion? I am a computational social scientist who has spent his entire career studying how social media shapes political polarization. Several years ago, I became so concerned about political tribalism that I founded the Polarization Lab—a team of social scientists, statisticians, and computer scientists at Duke University, where I am a professor. Our team diagnoses the problems with our platforms using scientific research and builds new technology to reverse the course. Together, my colleagues and I have collected hundreds of millions of data points that describe the behavior of thousands of social media users over multiple years. We’ve run new kinds of experiments with automated social media accounts, conducted some of the first studies of how foreign misinformation campaigns influence people, and ventured deep inside social media companies to help them fight polarization. We’ve even created our own social media platform for academic research—allowing us to turn on and off different features of platforms to identify better ways of connecting people.

## THE LEGEND OF THE ECHO CHAMBER

This work has led me to question the conventional wisdom about social media echo chambers, but it has also inspired me to ask much deeper questions. Why does everyone seem so extreme on social media? Why do people like Dave Kelly spend hours arguing with strangers, even when they don't think it will change anyone's mind? Is using social media a temporary addiction that we can shake—like smoking—or is it fundamentally reshaping who we are and what we think of each other? No amount of data science wizardry can answer these questions. Instead, I wanted to see social media through the eyes of the people who use it each day. This is why our lab spent hundreds of hours interviewing people like Dave Kelly and carefully reconstructing their daily lives on- and off-line. And it's why I'm going to tell you the story of a recently bereaved extremist who lives in a motel where he wakes up and falls asleep watching Fox News—and a moderate Democrat who is terrified about school shootings but worries that posting his views on social media might cost him his job. These stories not only help me paint a more complete picture of how political polarization unfolds on social media; they also inspired me and my colleagues to run new types of large-scale experiments in turn.

Studying social media from the perspective of the people who use it is also important because they are conspicuously absent from public debates about social media and political tribalism. Instead, our current conversation is dominated by a handful of tech entrepreneurs and software engineers who helped build our platforms. These Silicon Valley apostates now claim the technology they created wields unprecedented influence over human psychology—technology that not only traps us within echo chambers, but also influences what we buy, think, or even feel. Facebook, Twitter, and other platforms were either asleep at the wheel when malicious foreign actors launched campaigns to influence social media users—these apostates claim—or willfully ignored them because they increased user engagement (and

## CHAPTER 1

therefore their bottom line). This narrative is very seductive for anyone searching for a scapegoat for our current situation, but is it really true? Though social media companies are by no means blameless for our current situation, the evidence that people are simple dupes of political microtargeting, foreign influence campaigns, or content recommendation algorithms is surprisingly thin.

Instead, I will argue that our focus upon Silicon Valley obscures a much more unsettling truth: the root source of political tribalism on social media lies deep inside ourselves. We think of platforms like Facebook and Twitter as places where we can seek information or entertain ourselves for a few minutes. But in an era of growing social isolation, social media platforms have become one of the most important tools we use to understand ourselves—and each other. We are addicted to social media not because it provides us with flashy eye candy or endless distractions, but because it helps us do something we humans are hard-wired to do: present different versions of ourselves, observe what other people think of them, and revise our identities accordingly. But instead of a giant mirror that we can use to see our entire society, social media is more like a prism that refracts our identities—leaving us with a distorted understanding of each other, and ourselves. The social media prism fuels status-seeking extremists, mutes moderates who think there is little to be gained by discussing politics on social media, and leaves most of us with profound misgivings about those on the other side, and even the scope of polarization itself.

If social media platforms are so deleterious to democracy, why not delete our accounts? After all, I might enjoy using carrier pigeons to communicate my latest musings on Justin Bieber. But deleting our accounts is just not realistic. Social media has become so woven into the fabric of our lives—and particularly those of young people—that it is here to stay. The good news is this: if we social media users are the main source of political polar-

## THE LEGEND OF THE ECHO CHAMBER

ization, this means we also have the power to push back against it. In the chapters that follow, I'll describe how you can learn to see the social media prism and understand how it distorts the political landscape. I'll explain how we can begin to break the prism by changing our behavior and introduce you to new tools that my colleagues and I created in the Polarization Lab to help you do it. In addition to these "bottom-up" solutions, I offer a new path from the top down. I'll explain how social media platforms could be redesigned to bring us together, instead of pushing us apart. But first, I need to explain why breaking our echo chambers is the wrong place to start.

## A Better Social Media

**I**N MARCH 2020, the number of Google searches that included the term “unprecedented” soared higher than ever before. As the sheer scale of the COVID-19 pandemic set in, the watershed events that had occurred in the previous few years—the divisive 2016 election, two major investigations of a sitting president, and what could be described as the most dysfunctional period in the history of the U.S. Congress—felt like a distant memory. Only two months later, the virus had claimed more than sixty-five thousand American lives. And while the loss of lives was utterly devastating, the economic cost of the pandemic was also heart-stopping. After many companies closed or cut back their output, unemployment levels spiked near record highs. It became a cliché to say that the pandemic had stretched into every corner of American life, and hundreds of millions of Americans sheltered in place while they watched the grim news pour in, with each day feeling more unprecedented than the last.

I don’t think Muzafer Sherif, the temperamental social psychologist I profiled at the beginning of chapter 4, could have imagined a more perfect common enemy to force Democrats and Republicans to set aside their differences. Like the forest fire Sherif once planned to ignite to create solidarity between the summer campers he studied, COVID-19 seemed poised to burn everything to the ground. Not only did the pandemic devastate health and the economy (things that everyone cares about), but

saving both required unprecedented societal coordination in the form of social distancing. History teaches that crises—even those much less daunting than the COVID-19 pandemic—provide warring factions with the opportunity to hit the reset button.<sup>1</sup> To give one of many possible examples, the decade after the end of World War II witnessed unprecedented international cooperation and prosperity. Such crises also strengthen collective identities, both good and bad. In fact, some scholars argue that the Cold War between the United States and the Soviet Union kept U.S. political polarization in check for years.<sup>2</sup>

But there was something unique about the COVID-19 pandemic. During World War II, Republicans and Democrats huddled together in foxholes, where their shared fate was immediately apparent. But during the pandemic, Americans were intentionally avoiding as many other people as possible. According to data from Safegraph, a company that aggregates GPS data from smartphones for advertising purposes, the median number of minutes Americans spent inside their homes each day shot through the roof in March 2020.<sup>3</sup> The bars, restaurants, and other social venues that occasionally allowed Republicans and Democrats to interact were mostly shut down. By early May 2020, more than 100,000 small businesses had already closed their doors permanently, and many more soon followed.<sup>4</sup> Though some businesses eventually reopened, smartphone data from the summer of 2020 showed that many of their customers continued to stay home—perhaps because they were concerned about their health or that of others, or because they no longer had any disposable income.<sup>5</sup>

As Americans quarantined in their apartments, condos, and houses, the internet lit up. Remarkably, internet infrastructure survived the massive influx of users both new and old, who flocked online at (you guessed it) unprecedented rates. Within the first few weeks of social distancing, the amount of daily traffic to Netflix and YouTube increased by more than 15 percent.<sup>6</sup>

## CHAPTER 9

Social media usage increased even more. During the same time period, Facebook usage skyrocketed by 27 percent. Americans also flocked to new platforms. Zoom connected parents working from home with their colleagues, while they struggled to set up Google Classroom for their children—whose entire education had suddenly moved online. Use of Nextdoor.com, which connects people to their neighbors for discussions about community issues, increased by 73.3 percent. Houseparty, an app that connects friends for video chats during which they can play games together, saw a 79.4 percent increase in users.

How did the surge of social media activity shape political polarization on our platforms? Within days of the first reported cases of COVID-19 in the United States, pundits launched volatile debates on social media about closing the country's borders, the adequacy of testing regimes, and whether enough protective medical equipment was available.<sup>7</sup> These debates continued throughout the first two months of social distancing. But as some states announced plans to reopen their economies despite increasing infection rates within their borders, the discussion on social media grew more fervent. Liberal pundits warned that the second wave of infection would be worse than the first. Conservative pundits responded that the cure should not be worse than the disease, with some people worrying that suicides might outnumber COVID-19 deaths because of high unemployment. In my neighborhood, liberals shamed conservatives for allowing their children to play with others on Nextdoor.com. Facebook and Twitter were also quickly awash with misinformation about the coronavirus, ranging from claims that China created the disease in a lab to fabricated pictures of dolphins swimming in Venetian canals.

Meanwhile, nationally representative surveys of Americans painted a much different picture. Though Democrats expressed support for wide-scale protective measures against COVID-19 more quickly than Republicans did, a strong bipartisan consensus quickly emerged in favor of social distancing in early 2020.

## A BETTER SOCIAL MEDIA

Democratic leaders objected to the Trump administration's decision to close the country's borders in late February, yet by late March, 96 percent of Republicans and 94 percent of Democrats supported the move.<sup>8</sup> Though many Republican leaders objected to shelter-in-place orders from local governments, surveys revealed strong public support for these measures as well. In mid-March, a large survey revealed that 84 percent of Republicans and 89 percent of Democrats supported canceling gatherings of more than ten people.<sup>9</sup> The bipartisan consensus about the need for social distancing and other new measures persisted even as many governors announced plans to reopen their states' economies. A representative survey in late April of voters in Georgia, Tennessee, and Florida—three of the states that opened first—found that only 12 percent of voters supported reopening the economy at that time.<sup>10</sup> Another survey found that the number of Americans who described their country as “very divided” dropped from 62 percent in October 2018 to just 22 percent in April 2020.<sup>11</sup>

Unfortunately, there was very little sign of unity on social media. Instead, the social media prism was hard at work. The extremists we met in chapter 5 were busy playing the blame game. Jamie Laplace, the liberal medical assistant from Alabama whose story I told at the beginning of that chapter, tweeted that Fox News anchors who initially downplayed the seriousness of the crisis should be put in jail. Meanwhile, Ed Baker, the lonely conservative who falls asleep watching DirecTV each night in a Nebraska motel, was claiming that the Chinese government not only had started the pandemic, but was also scheming to disrupt the U.S. supply chain in its wake. What about the moderates we met in chapter 6? Sara Rendon, the moderate conservative worried about upsetting her liberal aunt, did not post a single tweet during the first two months of the pandemic. And Pete Jackman, the moderate liberal whose cousin received a death threat? He was busy tweeting about video games. Once again, the social media prism was refracting the most extreme people on the

political spectrum and muting moderates, two processes that reinforce each other, emboldening extremists to express increasingly radical beliefs and disenchanting moderates—many of whom will mistake such extremism as typical of the other side, as we've seen throughout this book.

## Social Media in the COVID-19 Era

By summer 2020, surveys showed that the fragile consensus among Americans about how to respond to the pandemic had already fallen apart.<sup>12</sup> Though I had hoped that the pandemic would encourage American social media users to set aside their differences and reevaluate the purpose of our platforms, it now seems that the social media prism is as powerful as ever. The crisis has not changed the rules of the game or the behavior of its most prominent players. On Facebook, Twitter, and Instagram, extremists continue to flex their partisan identities, treating nearly every interaction as an opportunity for their side to win or the other side to lose. Even worse, opportunities for Republicans and Democrats to humanize each other in off-line settings—which were already becoming rare before the pandemic—will only become less frequent as the virus changes how many people think about social interaction with strangers.

In chapter 8, I discussed how we can hack the social media prism to beat back polarization on our platforms. But I think that the only way we can create lasting improvement is to create a new playing field. Calling for a new form of social media may seem like a moon shot given Facebook's enormous market share. But taking the long view teaches us that platforms come and go. Friendster, the first broadly successful social media site, was once so popular that Google offered to buy it for \$30 billion.<sup>13</sup> Only two years later, MySpace supplanted Friendster, eventually becoming the most visited site on the internet by 2006.<sup>14</sup> And while Facebook once seemed unlikely to spread

## A BETTER SOCIAL MEDIA

beyond the elite college campuses where it was born, it would of course go on to exceed the expectations of Silicon Valley's most bullish prophets.

In more recent years, Facebook has also shown signs of vulnerability. It took five years for Facebook to knock MySpace off the top of the social media hierarchy. But only four years after that, Instagram rocketed onto the scene, attracting fifty million users in its first two years of existence.<sup>15</sup> Instagram usage surged so quickly that Facebook wisely acquired the platform in April 2012, allowing the photo-sharing site to maintain its distinctive identity. But Facebook and Instagram's combined monopoly on the selfie market would not go unchallenged for long: by December 2013, Snapchat users were sending more than 400 million messages per day, and by April 2015, the platform had more than 100 million users.<sup>16</sup> Similarly, TikTok came out of nowhere. It arrived in the United States in 2018 and surpassed 100 million users just two years later.<sup>17</sup> Though Facebook remains the dominant platform today, it continues to shed young people—whose preferences will ultimately determine the landscape of the market.<sup>18</sup>

What will come next? A new platform probably will not supplant the dominant platforms and serve as a forum for more productive discussions about politics. But just as social media users have splintered onto new platforms according to their interests, hobbies, and professional needs, I think that there is room for a new platform for political discussion. Would everyone use it? Of course not. But decades of social science research indicate that most people get their opinions about politics from friends, family members, or colleagues who proactively seek information about politics, regularly engage with others about such information, and care enough about issues to try to influence people in their social network who trust their opinions.<sup>19</sup> What would better social media look like for these people, particularly those who are appalled by the current state of polarization on our platforms?

## A New Kind of Platform

Imagine, for a moment, that we could peel back the layers of social media and isolate the ones that give the social media prism its power. When I first did this thought experiment several years ago, I had the urge to contact friends who work at social media platforms to see if they would be interested in conducting an experiment that could shed some light on this issue. After all, they were already experimenting with various parts of their platforms. Twitter apparently did extensive research before expanding the number of characters it allows users to write, staggering the rollout of this change to assess its impact.<sup>20</sup> Instagram did the same thing with its “like” button, allowing users in some countries to see the number of likes their posts received but not displaying this information to others. The company learned that this change might discourage status seeking or bullying on the platform. Facebook runs so many experiments on its platform that data scientists within the company were forced to build an entirely new type of software to compare dozens—or even hundreds—of experiments to each other. However, the vast majority of these experiments are designed to improve advertising and user retention. The type of experiments my colleagues and I in the Polarization Lab wanted to run were quite different.

We wanted to ask fundamental questions about the architecture of social media. More specifically, we wanted to explore interventions that could disrupt the feedback loop between identity, status seeking, and political polarization that I have described throughout this book. As I mentioned at the end of chapter 8, we hoped to create a place where people like Sara Rendon and Derek Hunter could have a productive conversation about race and policing. The first thing our lab was interested in exploring was what might happen if you could create a space where people can put their ideas before their identities. What would happen, we wondered, if you brought Republicans and Democrats together

## A BETTER SOCIAL MEDIA

to discuss politics on a new platform where they were completely anonymous? Would moving the game to a new playing field allow people to find common ground more easily?

There is, of course, a very dark side to anonymous conversations on the internet. Anonymous websites such as 4chan.org host some of the most disturbing content on the internet.<sup>21</sup> Similarly, many anonymous Twitter trolls make outlandish statements that they would probably never make in person. Other platforms such as Yik Yak—a defunct platform that allowed people to converse anonymously with other people nearby—were also accused of encouraging uncivil or abusive behavior.<sup>22</sup> At the same time, other social media forums thrive precisely because of anonymity. Though anonymity can minimize the social consequences of bad behavior, it also gives people the opportunity to explore alternative views outside the social relationships that often prevent them from doing so.<sup>23</sup> A prime example is Change My View, a Reddit community whose users are invited to post a statement about a controversial topic so that others can attempt to persuade them to reconsider their position. Users award each other “deltas” if they changed their minds about an issue. Change My View (which spawned another app called Ceasefire) is not limited to political topics, though many of the discussions on these forums included spirited—yet highly rational—debates about social policies among sizeable groups of people.<sup>24</sup>

Studying whether anonymous conversation increases or decreases polarization would simply carry too much risk for large companies like Facebook or Twitter. Besides, our goal was to take Republicans and Democrats out of their environment and create a new playing field. So we decided to create our own social media platform, one that would allow us to experiment with different features of platforms in a highly ethical manner. We wanted the platform to be polished enough to effectively simulate the experience of real users interacting with each other. We also wanted

it to be attractive enough that people would actually want to use it. We did not want the platform to have a name that would give away its purpose. If we called it “the anti-echo chamber,” for example, we worried that our platform might attract the presumably small group of people who are willing to change their views—a criticism that is often leveled against proponents of Reddit’s Change My View.

We decided to give our platform a generic name: DiscussIt. We advertised it as “a place where people anonymously discuss various topics.” We hired a graphic designer to produce a sleek logo, and we spent more than a year writing, testing, and debugging the software for the platform. In the end, we created a mobile app that paired two people so they could discuss issues anonymously on a chat platform. After installing DiscussIt, users were shown a series of images that explain how to use the platform and told that they would be assigned an issue to discuss with a chat partner. DiscussIt informed them that their own names would be replaced with pseudonyms. All users were then assigned to a discussion topic and told that they would be matched with another user as soon as possible. Once users were matched, the app assigned them an androgynous name such as Jamie or Casey and directed them to the main chat interface. Users could then discuss topics in real time or asynchronously. The app periodically nudged people if they became nonresponsive, and I led a staff of graduate students who served as a user support team to further simulate a real social media platform. We even had to create a limited liability company to give our app the bona fides it needed to pass muster at the Apple AppStore.

To test whether DiscussIt could successfully reduce polarization between Republicans and Democrats, we designed a field experiment.<sup>25</sup> In early January 2020, we hired a survey firm to recruit 1,200 Democrats and Republicans to complete a survey about their political views as well as their attitudes toward people on the other side. Our survey also included a battery of questions

about immigration and gun control, the two issues we asked people to have a conversation about on DiscussIt in this study. One day after respondents completed the survey, we randomly assigned some of them to a treatment condition in which they were invited to earn \$17 to “test a new social media platform for a few hours.” Those who consented were given an invite code and instructed to plug it into the app’s log-in screen. Unbeknownst to the study participants, our app used this code to pair them with another person in the study who was from the opposing political party.<sup>26</sup> One week later, we sent all study participants a follow-up survey that included many of the questions asked in the first survey.

The results of the experiment make me cautiously optimistic about the power of anonymity. People who used DiscussIt exhibited significantly lower levels of polarization after using it for just a short time. Many people expressed fewer negative attitudes toward the other party or subscribed less strongly to stereotypes about them.<sup>27</sup> Many others expressed more moderate views about the political issues they discussed or social policies designed to address them. A lot of people became less polarized on all of these measures. Most surprising to me, however, is that an overwhelming majority of people told us they enjoyed using our social media platform, even though they had no incentive to do so. In a survey we administered within our app after users completed their discussions, 89 percent of people indicated that they had enjoyed their conversation. These engagements were not uniformly positive, but I was further encouraged by the people who emailed our user support team to ask how they could continue using our app after the testing session ended. Several users even asked how much the app would cost when it is released to the public.

If the quantitative results of the experiment were encouraging, then reading the transcripts of people’s interactions was doubly so. People shared heartbreaking stories about friends and

## CHAPTER 9

loved ones who had taken their own lives with handguns, as well as terrifying stories about being the victims of crime. One person even recounted in harrowing detail how she had survived a mass shooting. Others connected over nonpolitical issues. Two middle-aged parents sympathized with each other about raising teenage children, and others commiserated about friends or family members who blocked them because of their political views. Not every conversation was so congenial. One chat became so toxic that we decided to shut it down before it concluded. However, the vast majority of chats were not only civil but also highly productive. Many people outlined sophisticated policy proposals, and one pair even laid the groundwork for a grassroots campaign to prevent suicides in rural communities and became Facebook friends to make sure that they could continue their conversation after their trial of DiscussIt ended.

Our experiment also gives me hope that anonymous forums might facilitate more productive conversations across racial lines. One of our chats paired a White man who is a Republican and a Black woman who is a Democrat to discuss gun control. The man began the conversation by sharing his anger about the murder of two police officers in his home state of California. The woman responded with sympathy, telling the man that her son is a police officer. But the woman went on to denounce recent incidents of police violence against unarmed Blacks as well. The conversation continued respectfully, and ultimately both participants moved toward slightly more moderate positions on handgun control, according to our follow-up survey. As I read this chat transcript, I began to wonder if this type of progress would have been possible if the participants had been aware of each other's racial background. Would the man have shared his views about the murder of police? Would the woman have had the opportunity to tell him that she shares his fear and anger? Would they have found any common ground at all? These questions are for future researchers to grapple with, but it's worth noting that most

## A BETTER SOCIAL MEDIA

of the 108 conversations in our study that connected Whites and Blacks had depolarizing effects that were on par with the other 660 conversations we analyzed. And race may only be the starting point. Could focused anonymous conversation also eliminate prejudice or status hierarchies associated with gender, class, or whether or not people have blue check marks next to their names (or other indicators of online celebrity)? People should not have to be anonymous to be heard or respected, but anonymity may give disadvantaged groups the power to decide how much they want to share about their identity. I'm hopeful that carefully designed anonymous conversations can help build the foundation for the deeper form of respect and mutual understanding that we need so badly.

### Platforms with a Purpose

The DiscussIt experiment is one of the few success stories in fighting political polarization on our platforms that I know of. But could a platform like DiscussIt really scale? Even if people were willing to jettison parts of their identities for more productive debates in the short term, would they continue to do so as the platform expanded to include thousands—or even millions—of users over weeks, months, or years? Would the dialogue remain so civil after trolls or extremists arrived? And would the promise of more-productive conversations about politics be compelling enough to make people keep coming back? Or would a platform like ours make an inglorious exit to the graveyard of social media?

To answer these questions, we need to revisit the core argument I've made throughout this book about why people use social media in the first place. In chapter 4, I argued that people keep coming back to social media because they help us do something that makes us distinctively human: create, revise, and maintain our identities to gain social status. Social media allow people to present different versions of themselves, monitor how

## CHAPTER 9

others react to those versions, and revise their identities with unprecedented speed and efficiency. But we humans are notoriously bad at judging what other people think of us—and the fleeting interactions we have with each other on social media make matters even worse. As I described in chapters 5 and 6, the social media prism fuels extremism, mutes moderates, and leaves most of us with profound misgivings about the other side. But we won't stop using social media any sooner than we will stop caring about our identities and social status. Instead, we need to think more about how the design of our platforms shapes the types of identities we create and the social status we seek.

What is the purpose of Facebook? The company tells us its mission is to “bring the world closer together.”<sup>28</sup> But the platform began as a sophomoric tool that Harvard undergraduates used to rate each other's physical attractiveness. What is the purpose of Twitter? Its motto is to “serve the public conversation,” but it was reportedly built to help groups of friends broadcast SMS-style messages to each other.<sup>29</sup> What is the purpose of Instagram? We're told it is to “capture and share the world's moments.” But the app was originally called “Burbn” (as in the drink) and was built to help people make plans to hang out with their friends.<sup>30</sup> What is the purpose of TikTok? I'm not even going to go there. Hopefully, my point is already clear: Should we really expect platforms that were originally designed for such sophomoric or banal purposes to seamlessly transform themselves to serve the public good? Should we be surprised when they create the kind of leaderless demagoguery from which anyone can invent a kind of status, no matter how superficial or deleterious to democracy? Is it any wonder that people find themselves so rudderless on social media, when there is no common purpose for posting in the void?

Imagine if we created a platform on which status was tied to a more noble purpose. Imagine a platform that gave people status not for clever takedowns of political opponents but for

## A BETTER SOCIAL MEDIA

producing content with bipartisan appeal. Once we articulate the purpose of our platforms more clearly, we can build such principles into the architecture of the entire system. Instead of boosting content that is controversial or divisive, such a platform could improve the rank of messages that resonate with different audiences simultaneously. Instead of recommending that people follow others who already share their views, a platform could expose them to those within their latitude of acceptance. We could define this range by asking people to complete a questionnaire about their views before they sign up for the platform. To fight false polarization, we could even build in the type of tools we have developed in the Polarization Lab. “Like” counters could be replaced by meters that show how people from across the ideological spectrum respond to people’s posts in blue, red, and purple. Artificial intelligence could be used to ask people who are about to post uncivil or ad hominem content to reflect upon their goals or to help people rephrase their messages using values that appeal to the other side.

Needless to say, not everyone would use a platform where you gain status for bridging political divides. But that may be a good thing. The trolls and extremists who gain notoriety on other platforms for taking down political opponents could not be entirely banned from the type of platform I’m envisioning, but it would be a much less rewarding place for them to play.<sup>31</sup> Instead of gaining attention for taking people down, their posts would be downranked because they only appeal to one side. Regulating extreme content would also become easier. Once platforms have a purpose, such as bridging partisan divides, it will be much easier to make policies to define acceptable and unacceptable behavior. Instead of moderation policies that are so broad that they are difficult or even impossible to enforce, a platform built to depolarize users could moderate any posts that were ad hominem or uncivil. These terms would still be difficult to define, but I think that a platform guided by top-down principles that are

## CHAPTER 9

transparent for all users to see when they sign up will always be more effective than a platform that relies upon norms to emerge from the bottom up—particularly when the bottom is composed of extremists whom moderates are unable or unwilling to police.<sup>32</sup> These policies—combined with undisclosed identity verification to further prevent trolling and extremism—could create the conditions necessary for anonymous conversation across party lines that is not only civil, but also productive.

Enforcing such policies would not be easy, or cheap. While the idea of a decentralized platform to reduce polarization from the bottom up is romantic, I don't think that it is realistic. Decentralized social media platforms such as Mastodon and Diaspora have struggled to attract new users while they try to keep the lights on.<sup>33</sup> A full-fledged effort to create a platform to depolarize us would require a significant investment from a major funder. Governments would be ideal candidates, and proponents of publicly funded platforms such as the internet activist Ethan Zuckerman have argued that they could be funded by a digital advertising tax.<sup>34</sup> Though depolarizing efforts should be an easy sell in highly polarized countries such as the United States, I will not hold my breath until they become a reality. If governments don't step in, there is money to be made by entrepreneurs willing to bet that a place where people can build reputations as effective bipartisan communicators will be more attractive to businesses, governments, and nonprofit organizations alike. There is already precedent for such a fundamental redesign of social media from companies such as Stack Overflow, where software developers gain reputational points for providing the best answers to each other's questions. Though this site originally served as a simple forum for the discussion of complicated technical issues, it has become so important to the technology industry that headhunters now sift through the platform to identify people who not only have technical talent but also enjoy the respect of their peers. As the United States confronts the most daunting challenge to

## A BETTER SOCIAL MEDIA

its social fabric in a generation, developing a reputation for bridging partisan divides may become an increasingly valuable social asset. And if the opportunities for Democrats and Republicans to engage with each other in off-line settings continue to decline, a new social media platform may be one of the few places where such reputations could be made.

It is also not too late for existing social media platforms to adopt many of the principles outlined above as well. Facebook, Twitter, and other platforms could optimize the order of posts in their users' timelines based on the amount of approval they generate from people on both sides of the aisle, instead of simple engagement metrics alone. Recommender algorithms could be trained to identify each user's latitude of acceptance—based on the content they like or who they follow—and encourage people to connect with people whose views are different, but not too different, from their own. Together, these two reforms could promote content that nudges social media users to engage with each other around issues where there is room for compromise, instead of getting worked up about the content that is enraging their own side on any given day. Our current platforms could also create new incentives for people to engage in such productive debate—for example, leaderboards that track how often prominent users generate content that appeals to people from both parties. Or, they could invent new forms of status—such as badges for people who effectively attract diverse audiences, instead of highlighting those who have achieved celebrity for their achievement outside of politics. Facebook, Twitter, and other dominant social media platforms could also experiment with creating new spaces on their platforms for people who want to reach across the aisle and offer their users the opportunity to connect using the same type of focused, small-scale conversations that we created on our simulated social media platform in the Polarization Lab.

Regardless of whether solutions come from the top down or the bottom up, the effort to depolarize our platforms could ben-

## CHAPTER 9

efit from a much deeper engagement with social science. As we unlock the keys to make our platforms less polarizing, we can use insights from social science to make them a reality. Instead of implementing untested interventions proposed by technology leaders, pundits, or policy makers, we must build the methods of empirical observation of human behavior into the architecture of our platforms, as some social media companies have already begun to do. Along the way, we must recognize that the immense challenges we face will continue to evolve over time. We must continue to test all of the solutions I've prescribed throughout this book in the coming years and decades. And if the rapidly evolving field of computational social science discovers the recommendations I've proposed no longer work, I'll be the first one calling for a change of course.