

CYBERCASCADDES

Any discussion of social fragmentation and online behavior requires an understanding of social cascades—above all because they become more likely when information, including false information, can be spread to hundreds, thousands, or even millions by the simple press of a button. Cascades are often hard or even impossible to predict, but they are all around us, and they organize our culture and even our lives. Increasingly, cascades are a product of social media. They occur within isolated communities, which develop a commitment to certain products, films, books, or ideas. Terrorists, rebels, and revolutionaries attempt to create and use them. Frequently cascades take hold far more generally, helping to produce (for example) a right to same-sex marriage, a rebellion against an authoritarian government, a nation's exit from the European Union, a new president, or a massively popular new cell phone.

It is obvious that many social groups, both large and small, move rapidly and dramatically in the direction of one or another set of beliefs or actions.¹ These sorts of cascades typically involve the spread of information; in fact, they are usually *driven* by information. Almost all of us lack direct or entirely reliable information about many matters of importance—whether George Washington actually lived, whether the earth goes around the sun, whether matter contains molecules, whether dinosaurs existed, whether there is a risk of war in India, whether the Islamic State of Iraq and the Levant (ISIL) is dangerous, whether a lot of sugar is actually bad for you, or whether Mars is real. For the vast majority of your beliefs, you really don't have direct information. You rely on the statements or actions of trusted others.

CYBERCASCADDES

TWO KINDS OF CASCADDES

To understand the social dynamics here, we need to distinguish between two kinds of cascades: informational and reputational.

Informational cascades. In an informational cascade, people cease relying at a certain point on their private information or opinions. They decide instead on the basis of the signals conveyed by others. It follows that the behavior of the first few people, or even one, can in theory produce similar behavior from countless followers.

To use a stylized example, suppose that Joan is unsure whether climate change is a serious problem. She may be moved to think that it is if her friend Mary thinks and says that climate change is a serious problem. If Joan and Mary are both favorably alarmed about climate change, their friend Carl may end up agreeing with them, at least if he lacks reliable independent information to the contrary. If Joan, Mary, and Carl believe that climate change is a serious problem, their friend Don will need to have a good deal of confidence to reject their shared conclusion. And if Joan, Mary, Carl, and Don present a united front on the issue, their other friends and even acquaintances may well go along. Something like this happens online every day.

It is important to emphasize a wrinkle here, which is that if one person sees that five, ten, a hundred, or a thousand people are inclined to say or do something, there is a tendency to think that each and every individual has made an independent decision to say or do it. The reality may well be that only a small fraction of the group made an independent decision. The rest are following the crowd, thus amplifying the very signal to which they were themselves subject. That signal may be extremely loud and seem quite impressive even though it incorporates the judgments of remarkably few people.

Environmental issues provide examples of how information travels, and can become quite widespread and entrenched, whether or not it is right. A disturbing illustration is the widespread popular belief that abandoned hazardous waste dumps rank among the

most serious environmental problems; science does not support that belief, which seems to have spread via cascade.² Another environmental example is the widespread and false belief that foods containing GMOs are hazardous to people's health; the scientific consensus is that they are not. Many cascades are widespread but local; consider the view, which had real currency in some African American communities in the 1980s, that white doctors are responsible for the spread of AIDS among African Americans. Or consider the notion, apparently widely held among American conservatives, that President Obama was not born in the United States—and the opinion, held by many parents and apparently defended at one point by Donald Trump, that vaccinations cause autism. One group may end up believing something, and another the exact opposite, and the reason is the rapid transmission of information within one group but not the other.

Even among specialists and indeed doctors, cascades are common. “Most doctors are not at the cutting edge of research; their inevitable reliance upon what colleagues have done and are doing leads to numerous surgical fads and treatment-caused illnesses.”³ Thus an article in the influential *New England Journal of Medicine* explores “bandwagon diseases” in which doctors act like “lemmings, episodically and with a blind infectious enthusiasm pushing certain diseases and treatments primarily because everyone else is doing the same.”⁴ It should be easy to see how cascades might develop among groups of citizens. And when informational cascades are operating, there is a serious social problem: people who are in the cascade do not disclose to their successors and the public the information (or reservations) that they privately hold.

Reputational cascades. We can also imagine the possibility of reputational cascades, parallel to their informational siblings.⁵ In a reputational cascade, people think that they know what is right, or what is likely to be right, but they nonetheless go along with the crowd in order to maintain the good opinion of others. Even the most confident people sometimes fall prey to this pressure, silencing themselves in the process. Fearing the wrath of others,

people might not publicly contest practices and values that they privately abhor.

The social practice of sexual harassment long predated the legal notion of “sexual harassment” and the innumerable women who were subject to harassment did not like it. But mostly they were silent, simply because they feared the consequences of public complaint. It is interesting to wonder how many current practices fall in the same general category—they produce harm, and are known to produce harm, but they persist because most of those who are harmed believe that they will suffer if they object in public. Whole governments can fall once reputational cascades start growing, as they often do when people learn that their disaffection is widely shared.

To see how a reputational cascade might work, suppose that Albert suggests vaccinations can cause autism, and Barbara concurs with Albert, not because she actually thinks that Albert is right, but because she does not wish to seem, to Albert, to be ignorant or indifferent to a serious risk faced by children. If Albert and Barbara seem to agree that vaccinations can cause autism, Cynthia might not contradict them publicly and might even appear to share their judgment, not because she believes that judgment to be correct, but because she does not want to face their hostility or lose their good opinion.

It is easy to see how this process might generate a reputational cascade. Once Albert, Barbara, and Cynthia offer a united front on the issue, their friend David might be most reluctant to contradict them even if he thinks that they are wrong. The apparent views of Albert, Barbara, and Cynthia carry information; that apparent view might be right. But even if David thinks that they are wrong and has information supporting that conclusion, he might be most reluctant to take them on publicly. Reputational cascades impose increasing pressure as larger numbers of people join the cascade. A position that was once highly unpopular, leading people to silence themselves, may come to seem widely held, so much so that people risk their reputations if they oppose it.

INFORMATION AS WILDFIRE AND TIPPING POINTS

The Internet greatly increases the likelihood of diverse but insistent cascades. Cybercascades occur every day. On Twitter and Facebook, you can find them in an instant. They might involve politics, miraculous products, deadly diseases, conspiracies, unsafe food, supposed events in Moscow or Berlin, or anything else.

Here is some fun and illuminating evidence of how online cascades can happen, from the domain of music.⁶ A team of experimenters, led by Matthew Salganik, Peter Dodds, and Duncan Watts, created an artificial music lab, including 14,341 participants. The participants were given a list of dozens of previously unknown songs from unknown bands; they were asked to listen to a brief selection of any songs of interest to them, decide what songs (if any) to download, and assign a rating to the songs they chose. About half the participants were asked to make their decisions independently, based on the names of the bands and the songs and their own judgment about the quality of the music. About half the participants could see how many times each song had been downloaded by other participants. These participants were also randomly assigned to one or another of eight possible “worlds,” or subgroups, with each evolving on its own; those in any particular world could see only the downloads in their own world. A key question was whether people would be affected by the choices of others—and whether different music would become popular in the different “worlds.”

Did social influences matter? Did cascades develop? There is not the slightest doubt. In all eight worlds, individuals were more likely to download songs that had been previously downloaded in significant numbers, and less likely to download those that had not been so popular. Most strikingly, the success of songs turned out to be almost entirely unpredictable! Almost all the songs could become popular or unpopular, with everything depending on the choices of the first downloaders. The identical song could be a hit or a failure—simply because other people, at the start, were seen

to choose to download it or not. (Think for a moment about how rumors spread or fail to spread on social media.)

To be sure, there is some relationship between quality and success. “In general, the ‘best’ songs never do very badly, and the ‘worst’ songs never do extremely well, but almost any other result is possible.”⁷ But even for the best and worst songs, there’s a high degree of unpredictability in terms of ultimate market shares, depending on whether they benefit from early popularity—and for the vast majority of songs, everything turns on social influences. Salganik, Dodds, and Watts acknowledge that in many ways, the real world is different from this experiment. They in fact controlled numerous variables, ensuring that their results are weaker than what happens in actual markets, where unpredictability is even greater, and where cascades are inevitable. Media attention, marketing efforts, critical reviews, and other pressures inflate the role of social influences. When experts fail to predict success, it is “because when individual decisions are subject to social influence, markets do not simply aggregate pre-existing individual preferences.”⁸ Note here that marketers often try hard to create early online “buzz” by suggesting that a certain cultural product is already popular; indeed, some marketing efforts actually involve artificial efforts to overstate the demand for the product, through purchases not by ordinary people, but by those allied with the artist.

Social media are full of such efforts. An acquaintance of mine, the author of an excellent book in the general domain of behavioral science, has tweeted on numerous occasions something like, “My book is doing great and well above expectations! Thanks for the support!” Actually the book isn’t doing so great, but the author knows well that if people think that other people are buying it, they’ll be more likely to buy it themselves.

Consider in this regard the 2013 Oscar winner for best documentary, *Searching for Sugar Man*, a stunning film about an unsuccessful Detroit singer-songwriter named Sixto Rodriguez, who released two long-forgotten albums in the early 1970s. Almost no one bought his albums, and his label dropped him. Rodriguez

stopped making records and worked as a demolition man. What Rodriguez didn't know, while working in demolition, was that he had become a spectacular success in South Africa—a giant, a legend, comparable to the Beatles and the Rolling Stones. Describing him as “the soundtrack to our lives,” South Africans bought hundreds of thousands of copies of his albums, starting in the 1970s. *Searching for Sugar Man* is about the contrast between the failed career of Detroit's obscure demolition man and the renown of South Africa's mysterious rock icon.

The film is easily taken as a real-world fairy tale, barely believable—a story so extraordinary that it gives new meaning to “you couldn't make it up.” But as the music lab experiment shows, it is a bit less extraordinary than it seems, and it offers a profound lesson not only for music and culture markets but for business and politics as well.

We like to think that intrinsic quality produces success, and that quality will ultimately prevail in free markets. To be sure, quality is usually necessary, but it's not enough. Social dynamics—who is conveying enthusiasm to whom, and how loudly, and where, and exactly when—can separate the rock icon from the demolition man, and mark the line between stunning success and crashing failure.

And if this is true for online music, it is likely to be so for many other things as well, including movies, books, political candidates, and even ideas. (“Everyone is flocking to candidate X,” or “idea Y is really catching on.”) Candidates and ideas may enjoy stunning success (or failure) simply because social dynamics give them an early boost (or not). Here we can see a large effect from collaborative filtering, which may help move or entrench, and not merely reflect, individual preferences.

POLITICAL CASCADES AND TURBULENCE

These points suggest a hypothesis: political life is a lot like the music lab. In fact, it is a kind of real-world politics lab. Bill Clinton, George W. Bush, Obama, and Trump succeeded not only because

of their evident talents but also because they received the equivalent of many early downloads. There are many talented politicians who never succeeded, and the reason isn't that they were not quite talented enough. It is that they failed to attract the right level of attention, either early or at some crucial time. The same is true for policy reforms.

It is not simple to test this hypothesis, but Helen Margetts, Peter John, Scott Hale, and Taha Yasseri have made significant strides in their book, *Political Turbulence*.⁹ Their subtitle is *How Social Media Shape Collective Action*, but their thesis is far more specific and striking. They argue that there is a great deal of unpredictability in modern political life, that the level of predictability is significantly increased by social media, and that social influences heighten unpredictability. Explicitly referring to the music lab experiment, they claim that in the age of social media, political movements are likely to be highly turbulent.

With respect to social influences, some of their best evidence comes from petitions. Both the United Kingdom and the United States have created online petition platforms. Most petitions fail—and fail quickly. No one pays them the slightest attention. As it turns out, the first day that a petition becomes public is critical. Early popularity makes all the difference, because political momentum builds on itself. In the United Kingdom, five hundred signatures are required to obtain an official response, and a large percentage of successful petitions get there within two days. It is reasonable to think that a certain (small) number of petitions spur and benefit from early cascade effects; they are a lot like Rodriguez in South Africa. But the vast majority of petitions fail to do that; they are just like Rodriguez in the United States.

That is indeed a reasonable thought, but it is not the only reading of the data. It is possible that some petitions receive large numbers of independent signatories, and that social influences do not much matter. But Margetts and her colleagues offer strong reasons to think otherwise. For one thing, social media have a large effect. The number of signatures and the number of tweets are closely

correlated; the more tweets, the more signatures. The authors' analysis of both timing and content suggests that tweets are driving signatures, rather than the other way around.

But the strongest evidence of the power of social influences comes from the fact that in April 2012, the UK Cabinet Office introduced "trending petitions" information on its web page so that everyone could see which petitions were succeeding and how many other people had signed. Margetts and her colleagues explore the effects of that information. Somewhat surprisingly, it had no effect on the overall level of petition signing. But it greatly affected the distribution of signatures. Using a method akin to that in the music lab study, the researchers find that *after the trending information was introduced, signatures were much more concentrated on a small number of petitions*. That is important evidence that "the information-rich get richer, and the information-poor get poorer."¹⁰ Note that we are speaking here of what kinds of petitions receive attention from high levels of government. On that question, what is observed mirrors the music lab experiment.

As one might expect, Margetts and her colleagues find that small design changes can have large and unintended consequences. The United Kingdom lists the top six petitions (measured by number of signatures) in order on its website, and it also provides visitors with the option to click to see six more. Margetts and her colleagues tested whether and how the trending information affected people's signatures. The details of the test need not detain us, but the central finding is major: the first-ranked positions received more concentrated attention—and signatures—as a result of that information. The upshot is that "the addition of the trending petitions facility causes the most popular trending petitions to receive more signatures, and that these signatures come at the expense of signatures to other petitions on the site."¹¹ We can undoubtedly reach broadly similar conclusions about how social media might promote or undermine political candidates. One result is unpredictability and turbulence, as modest differences in initial popularity map onto long-term variations.

As a further test of social influences, the researchers enlisted a website, WriteToThem, designed to help visitors write to public officials. The site reduces the costs of citizen engagement. In an experiment, people were randomly assigned to one of two groups. The first was the control, in which visitors to the site saw no social information. The second was the treatment, in which visitors could see how many others had written to a particular representative. Overall, 39 percent of those visitors who went to the page for their own representative ended up sending a letter. But there was a substantial difference between the two groups: 32.6 percent in the control, and 49.1 percent in the treatment.

Surprisingly, it did not matter, in the treatment group, whether the social information showed low, medium, or high levels of writing from previous visitors. What mattered was the information as such. One reason may be that people did not have the comparative information right before them; without a little work, they would not see that some representatives had higher percentages than others. Another reason may be that the differences between low (around 47 percent) and high (around 53 percent) were quite modest. With a wider range, we might expect something more like the music lab experiment and the petition data, where variations in numbers really did matter (and were visible).

This expectation is strongly supported by another experiment by the same researchers, testing people's willingness to sign petitions and pledge small donations on an assortment of political issues, such as climate change, protection of humpback whales, protection of the people of Darfur, negotiation of new trade rules to combat poverty, and protection of human rights in China. In the control group, people saw the petitions in random order. In the treatment group, people did so as well, but they were also given social information, stating whether there were already large numbers of signatures (over one million), small numbers (less than a hundred), or medium numbers (ranging from a hundred to one million).

Overall, people in the control group signed 61.5 percent of the petitions. As compared to the control, small and medium numbers

had no significant effect on whether people signed. But for those who saw large numbers, there was a real impact; that treatment group signed 66.7 percent of the time. True, that is not the level of effect that is observed in the music lab experiment (or the petition study). But it makes sense to think that the participants in this experiment were already inclined to sign, as reflected by the 61.5 percent overall signature rate, so that a massive difference among treatment conditions should not be expected. What matters is that high numbers had a consistent and statistically significant impact on the likelihood of signing—consistent in the sense that it cut across every one of the tested issues, notwithstanding varying levels of initial support.

Emphasizing that people with different personality traits show different propensities to engage, that it matters whether engagement is visible, and that people show different levels of susceptibility to social influences, Margetts and her colleagues conclude that “tiny acts,” made possible by social media, are “a growing form of political participation, which in some countries and contexts is overtaking voting as the political act that people are most likely to undertake.”¹²

Their most striking finding involves the nature of the underlying social dynamics. According to the researchers, “extroversion” predicts a willingness to participate at an early stage; if there is a significant number of extroverts, people with higher thresholds for participation might be moved—and once they are moved, those with lower thresholds will join, eventually encompassing large numbers of people. Because the costs of participating are so low (if only via a “like,” a retweet, or a signature), millions of people can form a movement in this way. And indeed, processes of this general kind seemed to have played a role in the collapse of authoritarian nations in North Africa—and in the fullness of time, they are likely to have large effects elsewhere as well.¹³

RUMORS AND TIPPING

On the Internet, rumors often spread rapidly, and cascades are frequently involved. Many of us have been deluged with e-mails about

the need to contact our representatives about some bill or other—only to learn that the bill did not exist, and the whole problem was a joke or a fraud. Even more of us have been earnestly warned about the need to take precautions against viruses that do not exist. In the 1990s, many thousands of hours of Internet time were spent on elaborating paranoid claims about alleged nefarious activities, including murder, on the part of President Clinton. Numerous sites, discussion groups, and social media posts spread rumors and conspiracy theories of various sorts. An old one: “Electrified by the Internet, suspicions about the crash of TWA Flight 800 were almost instantly transmuted into convictions that it was the result of friendly fire. . . . It was all linked to Whitewater. . . . Ideas become E-mail to be duplicated and duplicated again.”¹⁴ In 2000, an e-mail rumor specifically targeted at African Americans alleged that “No Fear” bumper stickers bearing the logo of the sportswear company of the same name really promote a racist organization headed by former Ku Klux Klan grand wizard David Duke.

Both terrorism and voting behavior have been prime areas for false rumors, fake news, and cascade effects. In 2002, a widely circulated e-mail said that a Boeing aircraft had not in fact hit the Pentagon on September 11. In 2004, many people were duly informed that electronic voting machines had been hacked, producing massive fraud. (If you’re interested in more examples, you might consult www.snopes.com, a website dedicated to widely disseminated falsehoods, many of them spread via the Internet.)

During the Obama presidency, countless e-mails were widely circulated about the alleged misconduct, incompetence, lying, disloyalty, and weirdness of President Obama and those who worked for him. The idea that Obama was born in Kenya (propagated by Donald Trump, among many others) is just one prominent example; another is that he is a Muslim. From 2009 to 2012, I had the honor of working in the Obama administration, and I was stunned to see the spread of false rumors about my own conduct and beliefs. (Some people said that I wanted to “steal people’s organs”; others said that I was behind Wikileaks.) What is especially interesting is

that those who believe such rumors need not be irrational. They are simply reacting to what other people seem to believe.

Most of these examples are innocuous, because no real harm is done, and because many cascades can be corrected. But as a disturbingly harmful illustration, consider widespread doubts in South Africa in the 1980s about the connection between HIV and AIDS. Because the AIDS virus infected a significant percentage of the adult population, any such doubts were especially troublesome. South African president Thabo Mbeki was a well-known Internet surfer, and he learned the views of the “denialists” after stumbling across one of their websites. The views of the denialists were and are not scientifically respectable—but to a nonspecialist, many of the claims on their (many) sites seemed plausible. At least for a period, President Mbeki both fell victim to a cybercascade and, through his public statements, helped to accelerate one—the point where many South Africans at serious risk were not convinced of an association between HIV and AIDS. It is highly likely that this cascade effect produced a number of unnecessary infections and deaths. It literally killed people.

Recall the existence of cascade effects among those who believe that childhood vaccinations are harmful and can in particular cause autism. If apparently reliable reports suggest that vaccinations cause autism, many parents will refuse them. That’s hardly innocuous. It can result in illness and death. In fact, the Internet is a breeding ground for false information about health and risk avoidance. It also provides reams of truth and makes it available to all. But every day, damaging falsehoods spread through informational cascades; consider the problem of fake news.

With respect to information in general, there is even a tipping point phenomenon, creating a potential for dramatic shifts in opinion. After being presented with new information, people typically have different “thresholds” for choosing to believe or do something new or different. As the more likely believers—that is, people with low thresholds—come to a certain belief or action, people with somewhat higher thresholds then join them, soon producing a

significant group in favor of the view in question. At that point, those with still higher thresholds may join, possibly to a point where a critical mass is reached, making large groups, societies, or even nations “tip.”¹⁵ The result of this process can be to produce cascade effects, as large groups of people end up believing something—whether or not that something is true or false—simply because other people in the relevant community seem to believe that it is true.

There is a great deal of experimental evidence of informational cascades, which are easy to induce in the laboratory; real-world phenomena also have a great deal to do with cascade effects.¹⁶ Consider, for example, going to college, smoking, participating in political protests, voting for third-party candidates, striking, recycling, filing lawsuits, using birth control, rioting, or even leaving bad dinner parties.¹⁷ In all these cases, people are greatly influenced by what others do. Often a tipping point will be reached. Sometimes we give an aura of inevitability to social developments, with the thought that deep cultural forces have led to (for instance) an increase in smoking, protesting, or a candidate’s success, when in fact social influences have produced an outcome that could easily have been avoided. Social media provide an obvious breeding ground for cascades, and as a result, thousands or even millions of people who consult sources of a particular kind will move in one or another direction, or even believe something that is quite false.

The good news is that the Internet, including social media, is easily enlisted to debunk false rumors as well as start them. Online, people can correct those rumors in a hurry. For this reason, most such rumors do no harm. But it remains true that the opportunity to spread apparently credible information to so many people can induce fear, error, and confusion in a way that threatens many social goals, including democratic ones. As we have seen, this danger takes on a particular form in a balkanized speech market as local cascades lead people in dramatically different directions. When this happens, correctives, even via the Internet, may work too slowly or not at all, simply because people are not listening to one another. Recall the (terrible) problem of the backfiring correction.

UP AND DOWN VOTES

We continue to learn more about how social influences work online. Lev Muchnik, a professor at the Hebrew University of Jerusalem, and his colleagues carried out an ingenious experiment on a particular website—one that displays a diverse array of stories and allows people to post comments, which can in turn be voted “up” or “down.”¹⁸ With respect to the posted comments, the website compiles an aggregate score, which comes from subtracting the number of down votes from the number of up votes. To study the effects of social influences, the researchers explored three conditions: “up-treated,” in which a comment, when it appeared, was automatically and artificially given an immediate up vote; “down-treated,” in which a comment, when it appeared, was automatically and artificially given an immediate down vote; and “control,” in which comments did not receive any artificial initial signal. Millions of site visitors were randomly assigned to one of the three conditions. The question was simple: What would be the ultimate effect of an initial up or down vote?

You might well think that after so many visitors (and hundreds of thousands of ratings), a single initial vote could not possibly matter. Some comments are good, and some comments are bad, and in the end, quality will win out. It’s a sensible idea, but if you thought it, you would be wrong. After seeing an initial up vote (and recall that it was entirely artificial), the next viewer became 32 percent more likely to give an up vote too. What’s more, this effect persisted over time. After a period of five months, a single positive initial vote artificially increased the mean rating of comments by a whopping 25 percent! It also significantly increased “turnout” (the total number of ratings).

With respect to negative votes, the picture was not at all symmetrical—an intriguing finding. True, the initial down vote did increase the likelihood that the first viewer would also give a down vote. But that effect was rapidly corrected. After a period of five months, the artificial down vote had zero effect on median ratings (although it did increase turnout). Muchnik and his colleagues

conclude that “whereas positive social influence accumulates, creating a tendency toward ratings bubbles, negative social influence is neutralized by crowd correction.”¹⁹ They think that their findings have implications for product recommendations, stock market predictions, and electoral polling. Maybe an initial positive reaction, or just a few such reactions, can have major effects on ultimate outcomes—a conclusion very much in line with Salganik, Dodds, and Watts’s evidence. But maybe negative reactions will get corrected pretty quickly.

It’s an interesting thought, but we should be careful before drawing large lessons from a single study, particularly when participants had no money on the line. It’s possible that negative reactions can have long-term effects on products, people, movements, and ideas. But there is no question that when groups move in the direction of one or more of these, it may not be because of their intrinsic merits but instead because of the functional equivalent of early up votes. (Politicians, including Barack Obama and Donald Trump, often succeed as a result.) There are lessons here about the extraordinary unpredictability of groups—and their frequent lack of wisdom. Of course Muchnik and his colleagues’ own study involved large groups. But the same thing can happen in small ones, sometimes even more dramatically, because an initial up vote—in favor of some plan, product, or verdict—has a large effect on others.

HOW MANY MURDERS?

Here’s a clean test of group wisdom and social influences. The median estimate of a large group is often amazingly accurate. But what happens if people in the group know what one another are saying? You might think that knowledge of this kind will help, but the picture is a lot more complicated.

Jan Lorenz, a researcher in Zurich, worked with several colleagues to learn what happens when people are asked to estimate certain values, such as the number of assaults, rapes, and murders in Switzerland.²⁰ They found that when people were informed about

the estimates of others, there was a significant reduction in the diversity of opinions, which tended to make the crowd less wise.²¹

There's another problem with the crowd, which is that because people hear about other estimates, they also become more confident. Notably, people in the study received monetary payments for getting the right answer, so their mistakes were really mistakes—not an effort to curry favor with others. The authors conclude that for decision makers, the advice given by a group “may be thoroughly misleading, because closely related, seemingly independent advice may pretend certainty despite substantial deviations from the correct solution.”²² There's a lesson there for the wisdom of crowds in online settings. Because people are interacting with one another, they might not be so wise.

SEGREGATION, MIGRATION, AND INTEGRATION

The Daily Me is not a lived reality, at least for most of us. Facebook, Twitter, Instagram, and Snapchat accounts can certainly spread diverse points of view, and many people use them in exactly that way. Facts and opinions on liberal sites often migrate to conservative sites, and vice versa. We have seen that even if opinions are clustering, society can benefit from the wide range of arguments that ultimately make their way to the general public. And for many of us, voluntary choices do not produce clustering.

But there is also evidence of an echo chamber effect, at least for some of us. For example, a 2009 study finds modest but clear evidence of such an effect.²³ Examining the behavior of 727 people over a six-week period, R. Kelly Garrett found that people are significantly more likely to click on information that reinforces their views, and somewhat less likely to expose themselves to information that contradicts those views. In her account, people seek support for their own positions, and they do so consistently. It follows that people “are more likely to be interested in reading a story that they expect to support their opinion, and they spend more time reading it. They are also marginally less likely to be interested in

stories containing opinion-challenging information, but they do not systematically avoid them.”²⁴ The fact that people spend more time with stories that support their views is worth underlining.

The echo chamber effect here is not large: while people prefer information that supports their convictions, they do not run from information that undermines them. In Garrett's words, “People's desire for opinion reinforcement is stronger than their aversion to opinion challenges.” Her conclusion is that people “do not seek to completely exclude other perspectives from their political universe, and there is little evidence that they will use the Internet to create echo chambers, devoid of other viewpoints, no matter how much control over their political information environment they are given.”²⁵ In short, her study finds an inclination to find like-minded sources, but importantly, they are hardly sealed.

At the same time, Garrett offers an ominous projection: “Polarized news outlets serving niche audiences, which are more economically feasible online where production costs are lower, are another threat. Faced with a choice between a news source that is almost exclusively supportive of their opinions and another that almost exclusively challenges those same opinions, news consumers seem likely to choose the former.”²⁶ Garrett has done a great deal of work on these issues, and it is broadly consistent with her central findings here and also signals the existence of that threat.²⁷

One of the most systematic treatments of these issues comes from economists Matthew Gentzkow and Jesse M. Shapiro, who compare ideological segregation online and offline.²⁸ To measure ideological segregation, they use an “isolation index,” which, in their words, is

equal to the average conservative exposure of conservatives minus the average conservative exposure of liberals. If conservatives visit only foxnews.com and liberals only visit nytimes.com, for example, the isolation index will be equal to 100 percentage points. But if both conservatives and liberals get all their news from cnn.com, the two groups will

have the same kind of exposures, and the isolation index will be equal to 0.²⁹

That's a useful measure of segregation. Using data sources from 2004 to 2009, Gentzkow and Shapiro find a clear difference between what conservatives and liberals see online. On the Internet, the average conservative's exposure to conservative news is 60.6 percent, while the average liberal's is 53.1 percent, producing an isolation index for the Internet of 7.5 percentage points. That's significant, but again, it's not huge. You could easily see it as modest. Gentzkow and Shapiro find that most people are not using the Internet to live in echo chambers. For example, a consumer who received news exclusively from foxnews.com would have a more conservative news diet than 99 percent of Internet news users, which suggests that the vast majority of people are clicking on sites that do not fit a narrow political profile.

For four reasons, however, their data should be taken with some grains of salt, at least as applied to my concerns here. First, Gentzkow and Shapiro also find that isolation for the Internet is higher than for broadcast television news (1.8 percentage points), cable television news (3.3), magazines (4.7), and local newspapers (4.8)—though lower than that of national newspapers (10.4). The fact that it is higher than four standard sources of information is hardly comforting. Second, they are speaking of aggregate behavior, and the aggregate masks the extent to which significant subpopulations are creating echo chambers. Third, their findings are now dated; it is possible that the degree of isolation on the Internet is increasing. Fourth, more recent work finds that the echo chamber effect is dramatically higher on social media.

An intriguing qualification of the Gentzkow and Shapiro findings focuses directly on the question of subpopulations. Andrew Guess studied individual-level media consumption data to explore online behavior.³⁰ Looking at both surveys and browsing history, he finds that the percentage of visits that involve news and information about politics is actually quite low—about 6.9 percent of

all visits. Most of the time, people do not go online to explore politics. More relevantly for present purposes, both Democrats and Republicans do not sort themselves into echo chambers but instead tend to cluster around sites that can be counted as central, such as MSN.com and AOL.com. In general, Democrats and Republicans do not look radically different in their online behavior. The most important qualification is that Republicans also visit conservative sites (Townhall, the Drudge Report, and Breitbart) that Democrats entirely ignore. Democrats also show a greater interest than do Republicans in certain liberal sites (the Huffington Post and the Daily Kos). But overall, Guess finds “a remarkable degree of balance in respondents’ overall media diets regardless of partisan affiliation,” because most people’s choices “cluster around the center of the ideological spectrum.”³¹

It follows that most people do not consume news in a partisan way. But some people definitely do, including a set of left-leaning Democrats and (more pronounced in Guess's data) a subgroup of Republicans who visit conservative sites but not liberal ones. It follows that a small group of people is driving traffic to the most partisan outlets. Consistent with this finding, Guess also finds that in the aftermath of disclosures in 2015 about Hillary Clinton's use of a private e-mail server, Republicans suddenly flocked to increased consumption of news and information from identifiably conservative sources. Guess concludes that “a scandal that naturally maps onto the political divide involving a well-known figure can immediately drive traffic to more partisan sources.”³² A reasonable conclusion is that while most people do not live in echo chambers, those who do may have disproportionate influence, because they are so engaged in politics.³³

HOMOPHILY ON TWITTER

I have been discussing online behavior in general. What about social media? It is tempting to offer these hypotheses about Twitter in particular, consistent with my general concerns here: *people's Twitter*

feeds consist largely of like-minded types. When people retweet, it is generally because they agree with what they are retweeting. Because people generate their own feeds, they create echo chambers. To be sure, some people are at pains to say that a retweet “is not an endorsement,” but most of the time you retweet something because you like it and because you want your followers to see it as well. It is true that you might follow people with whom you do not agree, because you want to learn from them or you’re interested in knowing what the other side is thinking. But in general, we might hypothesize that Twitter is creating many thousands of information cocoons.

We can go a bit further. In business and government as well as the nonprofit sector, people are aware of the power of social media, and they use Twitter to their advantage. They try hard to create networks that will foster the preferred information environment. They tweet to produce positive impressions of their ideas and products—magazines, movies, television shows, books, candidates, and ideologies—and they have an intuitive awareness of group polarization and cascade effects. They create echo chambers by design.

Is that true? Though the full story is complicated and continues to emerge, there is considerable evidence that it is.³⁴ On Twitter, research finds a great deal of homophily. An important overview from 2001 finds homophily in all sorts of social networks, involving race, ethnicity, age, religion, and education, and constituting “niches” of identifiable kinds.³⁵ Eight years later, Gueorgi Kossinets of Google and Duncan Watts of Yahoo! Research (and now at Microsoft) explored the *origins* of homophily; with particular emphasis on the role of individual choices and structures. Investigating actual behavior, they find that over many “generations,” a seemingly small and modest preference for similar others can “produce striking patterns of observed homophily.”³⁶

Building on the basic concept, Irai Himmelboim and his coauthors find a great deal of homophily on Twitter. Studying Twitter networks involving ten controversial political topics, they find that

“Twitter users are unlikely to be exposed to cross-ideological content from the clusters of users they followed, as these were usually politically homogeneous.”³⁷ To be sure, people create social ties on Twitter on the basis of many common interests, not merely politics. In the political domain, interest in content is largely confined to like-minded users. There is little cross-ideological communication, at least of a meaningful kind.

Similarly, M. D. Conover and his coauthors investigated networks of political communication on Twitter, including more than 250,000 tweets from more than 45,000 users during the six weeks in advance of the 2010 US congressional midterm elections. Studying retweets, they find a massive degree of ideological segregation, with “extremely limited connectivity between left- and right-leaning users.” The “retweet network,” as the authors call it, separates users into two homogeneous communities, corresponding to the political right and left.³⁸ The sheer level of clustering is quite remarkable, suggesting two different communications universes.

At the same time, there is a puzzle—a finding that cuts in the other direction. If you look at “mentions” as opposed to retweets (“mentions” are tweets that contain another Twitter user’s @username), you will find far less in the way of political segregation. On Twitter, conservatives do mention tweets from people with liberal views, and liberals do mention tweets from people with conservative views. The authors speculate that hashtags might be a big reason. If someone issues a tweet that says #SecondAmendment, a lot of people might be interested in it, especially if it is seen as having a neutral or mixed valence.

But the authors ultimately conclude that the interactions they find “are almost certainly not a panacea for the problem of political polarization.” The problem is that despite the large number of mentions, ideologically opposed users “very rarely share information from across the divide with other members of their community.” Hence “political segregation, as manifested in the topology of the

retweet network, persists in spite of substantial cross-ideological interaction.³⁹

A study of Twitter data from the 2012 election cuts in the same direction.⁴⁰ On November 5 of that year—the day before the election—economists Yosh Haberstram and Brian Knight downloaded information from 2.2 million Twitter users who had followed the Twitter handles of candidates for the House of Representatives. The researchers coded Twitter users as liberal or conservative “voters” based on the party affiliation of the candidates they followed (for example, those who followed more Republican candidates were considered conservative), and then confirmed those ideologies based on the type of news outlets the voters followed (for instance, liberals were much more likely to follow *Hartball* with Chris Matthews). From these politically engaged Twitter “voters,” Haberstram and Knight analyzed 90 million links to other Twitter users as well as 500,000 candidate retweets and mentions of candidates.

The researchers found that people were disproportionately exposed to like-minded tweets. Specifically, the researchers discovered that conservative exposure among conservatives was 77.6 percent, but just 37.2 percent among liberals, yielding an isolation index of 40.3 percentage points on Twitter. That’s far higher than the 7.5 percent found by Gentzkow and Shapiro for ideological segregation on the Internet. How could that be?

To reconcile their study with Gentzkow and Shapiro’s findings, the researchers focused on two factors: people who follow politicians might strongly prefer to link to like-minded individuals, and news consumption on Twitter in particular might affect ideological segregation. The researchers found that the isolation index was indeed significantly lower (21.7 percent) for Twitter users who followed candidates from more than one party (so-called moderates).⁴¹ It was also significantly lower (24.1 percent) for media consumption—that is, among users who followed media outlets (think Fox News or *New York Times*), the level of ideological segregation was still significant but less pronounced than it was for those

who do not follow such outlets. In Haberstram and Knight’s words, “[the] same Twitter users experience[d] lower segregation when consuming news from media outlets than when using Twitter as a social network” by linking to other voters.⁴² If these two factors are put together, the isolation index turns out to be just 6.7 percent, close to Gentzkow and Shapiro’s finding.

It follows that if you use Twitter to follow both media outlets and candidates from more than one party—that is, if you’re a “moderate”—then your ideological exposure will be slightly skewed, but not by much, and it will be comparable to ideological segregation on the Internet in general. But if you use Twitter primarily to follow candidates from just one political party (which many people do), and if you do not follow media outlets, then you’ll be exposed to dramatically different and far more limited viewpoints.

What does all this mean? For many users, Twitter is more ideologically segregated than radio, newspapers, and the Internet. Indeed, the researchers found that among the House candidate tweets that liberal voters saw, 90 percent came, on average, from Democrats; similarly, 90 percent of the candidate tweets that conservative voters saw came, on average, from Republicans.⁴³ (If the exposure had been random, these Twitter voters would have seen about half Democratic and half Republican tweets.) All this means that Twitter makes it easy for people inclined to hear like-minded viewpoints to do exactly that—and many people are following their inclinations.

Studying Republicans and Democrats in the United States, Elnor Colleoni and her coauthors find a great deal of political homophily, but with some intriguing differences between Democrats and Republicans.⁴⁴ In brief, Democrats in general show significantly higher levels of political homophily, but Republicans who follow official Republican accounts show higher levels of homophily than do Democrats.

A great deal remains to be learned about the differences between Democrats and Republicans, and how these change over

time. In some years, one or another party will be more inclined to isolate itself on social media, and these inclinations probably shift from one period to another. Among both Democrats and Republicans, there are almost certainly differences between moderates and extremists. It is reasonable to speculate that those who consider themselves on the left wing of the Democratic Party are more inclined to homophily than those who consider themselves to be merely somewhat left of center, and something similar might be true of right-wing members of the Republican Party as compared to those who are merely somewhat right of center. It would also be intriguing and perhaps important to learn about the role of demographic characteristics. In Twitter, how does homophily differ between men and women, young and old, well educated and poorly educated, rich and poor?

In the fullness of time, an entire book should be written on this topic. It will undoubtedly complicate and qualify the intuitive hypotheses with which this section began. But the complications and qualifications are highly likely to be consistent with the claim that homophily is commonplace on Twitter, and that when millions of people use it to find news and opinions, birds of a feather are flocking together.

FRIENDS AND FACEBOOK

In general. What about Facebook? A study by its own employees strongly suggests that to some extent, Facebook's users are indeed creating political echo chambers.⁴⁵ Investigating how 10.1 million Facebook users interacted with news, the study explored the effects of Facebook's own (earlier) algorithm, which does a degree of filtering, and also users' own choices. One of the signal virtues of this study is that it cleanly separates the consequences of the Facebook algorithm from those of people's decisions whether or not to click. The authors' own emphasis is on the effects of the latter, with the suggestion that "the power to expose oneself to perspectives from the other side in social media lies first and foremost with

individuals."⁴⁶ That suggestion is not inconsistent with their actual findings, but the full story is more interesting.

Facebook's algorithm matters. As the authors' evidence shows, the algorithm suppresses exposure to diverse content by 8 percent for self-identified liberals and 5 percent for self-identified conservatives. That means that the algorithm will filter out one in thirteen crosscutting stories that a liberal might see, and one in twenty such stories that a conservative might see. True, those numbers are not huge, but they do mean that people are seeing (modestly) fewer news items that they would disagree with, solely because of the effects of the algorithm. And it shows Facebook's potential power to alter our news consumption: if people are getting a lot of their news from Facebook, the algorithm will create a skew.

With respect to individual choices, there is an additional and larger effect: clicking behavior results in exposure to 6 percent less diverse content for liberals and 17 percent less diverse content for conservatives. In this respect, the authors find clear evidence of *confirmation bias*, a product of motivated reasoning: people are more likely to click on material that confirms their beliefs and avoid material that undermines them. The best way to understand the study is to take the algorithm and individual choices together. In the aggregate, there is a great deal of self-sorting on Facebook, resulting in a situation in which people are likely to be seeing items with which they agree.

It is also true that as Facebook's researchers note, "Individuals do not encounter information at random in offline environments nor on the Internet."⁴⁷ And it is not so easy to measure how much less ideologically diverse information people are seeing on Facebook compared to face-to-face interactions or without Facebook's algorithm. Still, the figures do raise questions about what Facebook and other social media companies can or should do to promote ideological diversity. As we have seen, Facebook decided in 2016 to change its algorithm to prioritize posts by friends and family members over those of news publishers like the *Wall Street Journal* or *Huffington Post*.⁴⁸ That means that what you see on Facebook will

depend more on who your friends are, what they share, and what you click on. The change is highly likely to increase the echo chamber effect.

“Spend time viewing.” Facebook itself has a distinctive view about how to think about this situation. As Mark Zuckerberg, cofounder, chair, and chief executive officer of Facebook, once remarked, “A squirrel dying in front of your house may be more relevant to your interests than people dying in Africa.”⁴⁹ A clear example of the company’s commitment to consumer sovereignty is an upbeat blog post written by two people at the company in 2016. Revealingly, the post is titled “More Articles You Want to Spend Time Viewing.” The authors announce, “We are adding another factor to News Feed ranking so that we will now predict how long you spend looking at an article in the Facebook mobile browser or an Instant Article after you have clicked through from News Feed.” Cheerfully, they suggest that “with this change, we can better understand which articles might be interesting to you based on how long you and others read them, so you’ll be more likely to see stories you’re interested in reading.”⁵⁰

Without the slightest trace of self-consciousness, they add that the most recent changes to Facebook’s algorithm are intended to provide users “more articles [they] want to spend time viewing”—instead of the broad array of stories that users might not have otherwise considered (and on which they might not spend a whole lot of time). As algorithms become more accurate in the future, the company’s capacity to prescreen posts for what users want to read will inevitably improve. In a way, that’s great—but in a way, it really isn’t.

Science and conspiracies. A series of studies of Facebook users provides strong evidence that at least in certain domains, echo chambers exist on Facebook, and they are created by confirmation bias.⁵¹ One of those studies, led by Michela Del Vicario of Italy’s Laboratory of Computational Social Science, explores the behavior of Facebook users from 2010 to 2014.⁵² A central goal of the study was to test whether users create the virtual equivalent of gated communities.

Del Vicario and her coauthors examined how Facebook users spread conspiracy theories (using thirty-two public web pages), science news (using thirty-five such pages), and “trolls,” which intentionally spread false information (using two web pages). Their data set is massive; it covers all Facebook posts during the five-year period. The researchers looked at which Facebook users linked to one or more of the sixty-nine web pages, and whether they learned about those links from their Facebook friends.

In sum, the researchers find communities of like-minded people. Conspiracy theories, even if they are baseless, spread rapidly within such communities. On these issues, Facebook users tend to choose and share stories containing messages they accept—and neglect those they reject. If a story fits with what people already believe, they are far more likely to be interested in it and thus spread it. As Del Vicario and her coauthors put it, “Users mostly tend to select and share content according to a specific narrative and to ignore the rest.” On Facebook, the result is the formation of a lot of “homogeneous, polarized clusters.”⁵³ Within those clusters, new information moves quickly among friends (often in just a few hours).

The consequence is the “proliferation of biased narratives formulated by unsubstantiated rumors, mistrust, and paranoia.”⁵⁴ In that sense, confirmation bias is self-reinforcing, producing a vicious spiral. If people begin with a certain belief and find information that confirms it, they will intensify their commitment to that belief, strengthening their bias. Strong support for this conclusion comes from research from the same academic team, which finds that on Facebook, efforts to debunk false beliefs are typically ignored—and when people pay attention to them, they often strengthen their commitment to the debunked beliefs. The United States saw a lot of this during the 2016 presidential campaign.

Or consider how people respond to *intentionally* false claims. The researchers studied clearly unrealistic and satirical claims—for example, a post declaring that chemical analysis revealed that chemtrails contain sildenafil citratum (the active ingredient in

Viagra™).⁵⁵ The central finding is that many people liked and commented favorably on such claims. Even when information is deliberately false and framed with a satirical purpose, its conformity with the conspiracy narrative transformed it into suitable (and welcome) content for the relevant groups. To be sure, conspiracy theories, and those who like and spread them, are not exactly typical fare. We might expect to see an especially large echo chamber effect for such theories. But there is good reason to think that less extreme versions of the same general patterns of self-sorting can be found on Facebook.

Findings of this kind are important because people increasingly rely on social media for news. According to public opinion polls by the Pew Research Center, as of 2016, six out of ten US adults (62 percent) get news from social media, and 18 percent do so frequently. The polls also show that a majority of Twitter (59 percent) and Facebook users (66 percent) receive news on those platforms (both up significantly from 2013, when only about half these users got news there). Polls also demonstrate that while the percentage of the population that uses Twitter is relatively low (16 percent), Facebook is widely used (67 percent)—which means that some 44 percent of all US adults receive news from Facebook.⁵⁶

For people born after 1980, often called millennials, Facebook is by far the most common source of news about politics and government. In 2016, six out of ten millennials (61 percent) reported getting political news from Facebook, whereas only about four in ten (44 percent) said CNN, the next most popular source.⁵⁷ Facebook accounts for more than 40 percent of the referral traffic to news sites.⁵⁸ For better or worse, social media and Facebook in particular have a large effect in determining what people learn about political issues.

FACTS, VALUES, AND GOOD NEWS

To paraphrase an observation attributed to the late senator Daniel P. Moynihan, people are entitled to their own opinions, not to their own facts. But on some of the most politically charged issues,

people's ideological commitments settle their judgments about questions of fact. This point helps illuminate the effects of a fragmented media market; it contributes to polarization.

While many of the issues that divide people boil down to ideology and preference, there is at least one on which hard science should have a strong say—climate change. But do numbers and figures change people's opinions? In an experiment in 2016, my colleagues Sebastian Bobadilla-Suarez, Stephanie Lazzaro, Tali Sharot, and I asked more than three hundred Americans several climate-related questions, such as whether they believed that man-made climate change was occurring and whether the United States was right to support the recent Paris agreement to reduce greenhouse gas emissions.⁵⁹ On the basis of their answers, we divided participants into three groups: strong believers in man-made climate change, moderate believers, and weak believers.

Next we informed participants that many scientists have said that by the year 2100, the average temperature in the United States will rise at least 6 degrees Fahrenheit, and asked them for their own estimates of likely temperature rise by 2100. The overall average was 5.6 degrees Fahrenheit. As expected, there were significant differences among the three groups: 6.3 degrees for strong believers in man-made climate change, 5.9 degrees for moderate believers, and 3.6 degrees for weak believers.

Then came the important part of the experiment. Participants were randomly assigned to one of two conditions. Half of them received information that was more encouraging than what they originally received (good news for the planet and humanity); half of them received information that was less encouraging (bad news for the planet and humanity). In the good news condition, they were told to assume that in recent weeks, prominent scientists had reassessed the science and concluded the situation was far better than previously thought, suggesting a likely temperature increase of only 1 to 5 degrees. In the bad news condition, participants were told to assume that in recent weeks, prominent scientists had reassessed the science and concluded the situation was far worse than

previously thought, suggesting a likely temperature increase of 7 to 11 degrees. All participants were then asked to provide their personal estimates. Note that our experiment fits nicely with what happens online and in social media. All the time, people receive news, with respect to climate change, that suggests that the problem will be much better or much worse than previously thought.

Here's what we found. Weak believers in man-made climate change were moved by the good news; their average estimate fell by about 1 degree. But their belief was entirely unchanged by the bad news; their average estimate stayed essentially constant. By contrast, strong believers in man-made climate change were far more moved by the bad news (their average estimate jumped by nearly 2 degrees), whereas with good news, it fell by less than half of that (0.9 degrees). Moderate climate change believers were equally moved in both cases (they changed their estimates by approximately 1.5 degrees in each case).

The clear implication is that for weak believers in man-made climate change, comforting news will have a big impact, and alarming news won't. Strong believers will show the opposite pattern. As the media, including social media, expose people to new and competing claims about the latest scientific evidence, these opposing tendencies will predictably create political polarization, and it will grow over time.

There is a more general psychological finding in the background here. In the case of information about ourselves (about how attractive others perceive us to be, or how likely we are to get sick or to succeed), people normally alter their beliefs more in response to good news than in response to bad news. If you hear that you are better looking than you think, you will probably learn from that nice information. If you hear that you are not quite so good looking, you might well dismiss that unpleasant news. In certain circumstances, something similar will be true for political issues, as in the case of weak climate change believers, who are most likely to credit information suggesting that things will not be so bad. But at times, good political news can threaten our deepest

commitments, and we will be inclined give it less weight. Above all, we might want those commitments to be affirmed. Those who are most alarmed about climate change might prefer to learn that humanity really is at very serious risk than to learn that the climate change problem is probably not so bad. For them, bad news for humanity and the planet is, in a sense, good news (because it is affirming), and good news for humanity and the planet is, in an important sense, taken as bad news.

These findings help explain polarization on many issues, and the role of social media in increasing it. With respect to the Affordable Care Act, for example, people encounter good news, to the effect that it has helped millions of people obtain health insurance, and also bad news, to the effect that health care costs and insurance premiums continue to increase. For the act's supporters, the good news will have far more impact than the bad; for the opponents, the opposite is true. As the sheer volume of information increases, polarization will be heightened as well. Essentially the same tale can be told with respect to immigration, terrorism, and increases in the minimum wage. Which kind of news will have a large impact will depend partly on people's motivations and initial convictions.

But there's an important qualification. In our experiment, a strong majority showed movement; few people were impervious to new information. Most people were willing to change their views, at least to some extent. For those who believe in learning, and the possibility of democratic self-government, that's some good news.

IDENTITY AND CULTURE

A revealing body of research, coming largely from Yale Law School professor Dan Kahan, finds that "cultural cognition" shapes our reactions to science—and that our values affect our assessment of purely factual claims, even in highly technical areas.⁶⁰ As a result, Americans predictably polarize on factual questions involving, for example, gun control, climate change, nuclear waste disposal, and nanotechnology. Kahan's striking claim is that people's judgments

stem, in large part, from their sense of identity—of what kind of person they consider themselves to be. As a result, seemingly disparate views cluster. Among conservatives, for instance, gun control is a bad idea, and so is affirmative action; climate change is not a big problem; the Supreme Court should not have recognized same-sex marriage; and the minimum wage should not be increased.

In principle, it might be possible to identify specific values that link these apparently diverse conclusions. But Kahan's claim is that the real source of people's views, at least on certain controversial questions, is their understanding of their identity, and their effort to protect it. And while Kahan does not focus on online behavior and social media, there is no question that online interactions contribute to the phenomena he is describing.

Consider current debates over GMOs and climate change. The strong majority of scientists accept two propositions. First, GMOs generally do not pose serious threats to human health or the environment. Second, greenhouse gases are producing climate change, which does pose serious threats to human health and the environment.

With respect to GMOs, Democrats are far more likely than Republicans to reject the prevailing scientific judgment. With regard to climate change, Republicans are far more likely than Democrats to reject the prevailing scientific judgment. The partisan divide is easy to demonstrate. Among national leaders, many Democrats are concerned about GMOs; relatively few Republican leaders share that concern. Among ordinary citizens, a strong majority of Democratic voters believes that GMOs are unsafe. Republican voters are evenly divided on the safety question—a higher level of concern than that of their elected representatives, but much lower than that of Democratic voters.

In Congress, it is not exactly news that Democrats are far more likely than Republicans to support action to reduce greenhouse gas emissions. Take just one example: in a 2013 Senate vote on a non-binding resolution calling for a “fee on carbon pollution,” Republicans were in unanimous opposition, while most Democrats were

supportive. In recent years, about 75 percent of Democratic voters have said that they worry “a great deal” or “a fair amount” about climate change. For Republican voters, the percentage has ranged from 30 to 40 percent. As Kahan shows, Republicans who doubt climate change, and do not worry about it, do *not* display lower levels of scientific literacy.⁶¹ They are fully aware of what most scientists think. They are hardly ignorant. Their judgments appear to be a product of their values or sense of identity.

What is the best explanation for the fact that Republicans are more inclined to follow scientific opinion for GMOs, while Democrats are more inclined to do so for greenhouse gases? There are three contributing factors. Interest groups are the first. On the Democratic side, the concerns are of course sincerely held, but well-organized groups have been lobbying hard against GMOs, and they have been able to intensify public objections. These groups, which include the organic food industry and Whole Foods Market, have influence and credibility within the Democratic Party, and have stood to gain from mandatory labeling (which would harm their competitors). With respect to climate change, by contrast, the most powerful economic interests (such as the coal industry) have far greater influence within the Republican Party. Environmental groups, pressing for control of greenhouse gases, carry weight mostly with Democrats.

A second explanation points to my principal concern here: the effects of echo chambers, including social media. With respect to GMOs, some Democrats listen largely to one another, and their fears have become amplified as a result of internal discussion, even if science is not on their side. For greenhouse gases, the same phenomenon is occurring among Republicans. Here as elsewhere, discussions among like-minded people increase confidence, extremism, and polarization.

A third explanation builds on Kahan's research. It points to the crucial role of preexisting ideological commitments, which on particular issues can crowd out the effects of scientific findings. Many Republicans are opposed, in principle, to government interference

with free markets. They are inclined to be suspicious of scientific evidence that purports to justify that interference, especially in the environmental domain. By contrast, many Democrats are willing to indulge the assumption that corporate efforts to interfere with nature are potentially dangerous, especially if those efforts involve chemicals, new technologies, or pollution. Among Democrats, scientific claims about the risks associated with GMOs and greenhouse gases fall on receptive ears. In both cases, it is a matter of values first and scientific judgments second. And of course, a fragmented media market fortifies the relevant values.

To be sure, values do not always crowd out science. Some scientific questions do not trigger a sense of political identity; consider the question whether cigarettes cause lung cancer, or texting while driving increases the likelihood of accidents. Some scientific questions *migrate*: what was once a technical issue becomes politically inflamed, and what was once politically inflamed becomes technical.

As an example of the latter phenomenon, consider the depletion of the ozone layer, where the scientific evidence has long been overwhelming. That evidence led to bipartisan support for the Montreal Protocol, signed by President Ronald Reagan in 1988. Even so, there is no question that preexisting values help to account for political polarization with respect to GMOs and greenhouse gases. Taken together with the activities of interest groups and the echo chamber effect, those values help explain why the leaders of our two major political parties are strongly inclined to accept the dominant view within the scientific community in one case—and reject it in another.

The most unfortunate part is that interest groups, echo chambers, and conceptions of identity reinforce each other, creating a new kind of iron triangle. Interest groups use social media to promote their preferred view of the world as well as create or fortify conceptions of identity. The echo chambers increase the authority of those groups at the same time that they entrench those conceptions.

A CONTRAST: THE DELIBERATIVE OPINION POLL

By way of contrast to polarization and cybercascades, consider some work by James Fishkin, a creative political scientist at Stanford University who has pioneered a genuine social innovation: the deliberative opinion poll.⁶² The basic idea is to ensure that polls are not mere “snapshots” of public opinion. People’s views instead are recorded only after diverse citizens, with different points of view, have actually been brought together in order to discuss topics with one another.

Deliberative opinion polls have now been conducted in many nations, including the United States, England, and Australia. It is easy for deliberative opinion polls to be conducted on the Internet, and Fishkin has initiated illuminating experiments in this direction.

In deliberative opinion polls, Fishkin finds some noteworthy shifts in individual views. But he does not find a systematic tendency toward polarization.⁶³ In England, for example, deliberation led to a reduced interest in using imprisonment as a tool for combating crime.⁶⁴ The percentage believing that “sending more offenders to prison” is an effective way to prevent crime fell from 57 to 38 percent; the percentage believing that fewer people should be sent to prison increased from 29 to 44 percent; and belief in the effectiveness of “stiffer sentences” decreased from 78 to 65 percent.⁶⁵ Similar shifts were shown in the direction of greater enthusiasm for the procedural rights of defendants and increased willingness to explore alternatives to prison.

In other experiments with the deliberative opinion poll, shifts included a mixture of findings, with deliberation leading larger percentages of people to conclude that legal pressures should be increased on fathers for child support (from 70 to 85 percent), and that welfare and health care should be turned over to the states (from 56 to 66 percent).⁶⁶ These findings are broadly consistent with the prediction of group polarization, and to be sure, the effect of deliberation was sometimes to create an increase in the intensity with which people held their preexisting convictions.⁶⁷ But this was hardly a uniform pattern. On some questions, deliberation shifted a

minority position to a majority position (with, for example, a jump from 36 to 57 percent of people favoring policies making divorce “harder to get”), and it follows that sometimes majorities became minorities.⁶⁸

Fishkin’s experiments have some distinctive features. They involve not like-minded people but instead diverse groups of citizens engaged in discussion after being presented with various sides of social issues by appointed moderators. Fishkin’s deliberators do not seek to obtain a group consensus; they listen and exchange ideas without being asked to come into agreement. In many ways these discussions provide a model for civic deliberation, complete with reason giving.

It can be expensive, of course, to transport diverse people to the same place. But communications technologies make widespread uses of deliberative opinion polls as well as reasoned discussion among heterogeneous people far more feasible—even if private individuals, in their private capacity, would rarely choose to create deliberating institutions on their own. I have noted that Fishkin has created deliberative opinion polls on the Internet; there are many efforts and experiments in this general vein.⁶⁹ The social media can easily be enlisted for those purposes.

Here we can find considerable promise for the future in the form of discussions among diverse people who exchange reasons, and who would not, without current technologies, be able to talk with one another at all. If we are guided by the notion of consumer sovereignty, and if we celebrate unlimited filtering, we will be unable to see why the discussions in the deliberative opinion poll are a great improvement over much of what is now happening online. In short, aspirations for deliberative democracy sharply diverge from the ideal of consumer sovereignty—that is, a future in which, in Gates’s words, “you’ll be able to just say what you’re interested in, and have the screen help you pick out a video that you care about.”

But let’s offer a cautionary note: for many political questions, what matters is getting the facts straight, and for that, you need experts, not deliberative opinion polls. Suppose that the question is

whether current levels of particulate matter (an air pollutant) cause two deaths annually or two hundred, or instead two thousand. Or suppose the question is whether a requirement for greater fuel economy in trucks would produce less safe vehicles. On such questions, expertise is crucial. True, we might want the experts to deliberate. But a deliberative opinion poll might lead us in the wrong direction, even if people get pretty well informed.

Nonetheless, deliberative opinion polls are a lot better than non-deliberative ones. An enduring question is what sort of ideals we want to animate our choices, and what kinds of attitudes and regulations we want in light of that judgment. And here it is important to emphasize that current technologies, including social media, are in themselves hardly biased in favor of homogeneity and deliberation among like-minded people. Everything depends on what people seek to do with the new opportunities that they have. Consider the reflections of one Internet entrepreneur: “I’ve been in chat rooms where I’ve observed, for the first time in my life, African-Americans and white supremacists talking to each other. . . . [I]f you go through the threads of the conversation, by the end you’ll find there’s less animosity than there was at the beginning. It’s not pretty sometimes. . . . [b]ut here they are online, actually talking to each other.”⁷⁰ The problem is that this is far from a universal practice.

OF DANGERS AND SOLUTIONS

I hope that I have shown enough to demonstrate that for citizens of a heterogeneous democracy, a fragmented communications market creates a considerable number of dangers. There are dangers for each of us as individuals; constant exposure to one set of views is likely to lead to errors and confusions, sometimes as a result of cybercascades. And to the extent that the process entrenches existing views, spreads falsehood, promotes extremism, and makes people less able to work cooperatively on shared problems, there are dangers for society as a whole.

To emphasize these dangers, it is unnecessary to claim that people do or will receive all their information online. There are many sources of information, and some of them will undoubtedly counteract the risks I have discussed. Nor is it necessary to predict that most people are speaking only with those who are like-minded. Of course many people seek out or otherwise encounter competing views. But when technology makes it easy for people to wall themselves off from others, there are serious risks for the people involved and society as a whole.

5

SOCIAL GLUE AND SPREADING INFORMATION

Some people believe that freedom of speech is a luxury. In their view, poor nations, or nations struggling with social and economic problems, should be trying not to promote democracy but instead to ensure material well-being—economic growth, and a chance for everyone to have food, clothing, and shelter. This view is badly misconceived. If we understand what is wrong with it, we will have a much better sense of the social role of communications.

For many countries, the most devastating problem of all consists of famines, defined as the widespread denial of access to food and, as a result, mass starvation. In China's famine of the late 1950s, for example, about thirty million people died. Is free speech a luxury for nations concerned about famine prevention? Would it be better for such nations to give a high priority not to democracy and free speech but instead to economic development? Actually these are foolish questions. Consider the remarkable finding by the economist Amartya Sen that in the history of the world, there has *never* been a famine in a system with a democratic press and free elections.¹ Sen's starting point, which he also demonstrates empirically, is that famines are a social product, not an inevitable product of scarcity of food. Whether there will be a famine as opposed to a mere shortage depends on people's "entitlements"—that is, what they are able to get. Even when food is limited, entitlements can be allocated in such a way as to ensure that no one will starve.

entitlements
288