WEAPONS OF MATH DESTRUCTION

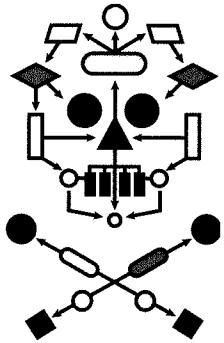

HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY

CATHY O'NEIL

CROWN
NEW YORK

INTRODUCTION

When I was a little girl, I used to gaze at the traffic out the car window and study the numbers on license plates. I would reduce each one to its basic elements—the prime numbers that made it up. $45 = 3 \times 3 \times 5$. That's called factoring, and it was my favorite investigative pastime. As a budding math nerd, I was especially intrigued by the primes.

My love for math eventually became a passion. I went to math camp when I was fourteen and came home clutching a Rubik's Cube to my chest. Math provided a neat refuge from the messiness of the real world. It marched forward, its field of knowledge expanding relentlessly, proof by proof. And I could add to it. I majored in math in college and went on to get my PhD. My thesis was on algebraic number theory, a field with roots in all that

factoring I did as a child. Eventually, I became a tenure-track professor at Barnard, which had a combined math department with Columbia University.

And then I made a big change. I quit my job and went to work as a quant for D. E. Shaw, a leading hedge fund. In leaving academia for finance, I carried mathematics from abstract theory into practice. The operations we performed on numbers translated into trillions of dollars sloshing from one account to another. At first I was excited and amazed by working in this new laboratory, the global economy. But in the autumn of 2008, after I'd been there for a bit more than a year, it came crashing down.

The crash made it all too clear that mathematics, once my refuge, was not only deeply entangled in the world's problems but also fueling many of them. The housing crisis, the collapse of major financial institutions, the rise of unemployment—all had been aided and abetted by mathematicians wielding magic formulas. What's more, thanks to the extraordinary powers that I loved so much, math was able to combine with technology to multiply the chaos and misfortune, adding efficiency and scale to systems that I now recognized as flawed.

If we had been clear-headed, we all would have taken a step back at this point to figure out how math had been misused and how we could prevent a similar catastrophe in the future. But instead, in the wake of the crisis, new mathematical techniques were hotter than ever, and expanding into still more domains. They churned 24/7 through petabytes of information, much of it scraped from social media or e-commerce websites. And increasingly they focused not on the movements of global financial markets but on human beings, on us. Mathematicians and statisticians were studying our desires, movements, and spending power. They were predicting our trustworthiness and calculating our potential as students, workers, lovers, criminals.

This was the Big Data economy, and it promised spectacular gains. A computer program could speed through thousands of résumés or loan applications in a second or two and sort them into neat lists, with the most promising candidates on top. This not only saved time but also was marketed as fair and objective. After all, it didn't involve prejudiced humans digging through reams of paper, just machines processing cold numbers. By 2010 or so, mathematics was asserting itself as never before in human affairs, and the public largely welcomed it.

Yet I saw trouble. The math-powered applications powering the data economy were based on choices made by fallible human beings. Some of these choices were no doubt made with the best intentions. Nevertheless, many of these models encoded human prejudice, misunderstanding, and bias into the software systems that increasingly managed our lives. Like gods, these mathematical models were opaque, their workings invisible to all but the highest priests in their domain: mathematicians and computer scientists. Their verdicts, even when wrong or harmful, were beyond dispute or appeal. And they tended to punish the poor and the oppressed in our society, while making the rich richer.

I came up with a name for these harmful kinds of models: Weapons of Math Destruction, or WMDs for short. I'll walk you through an example, pointing out its destructive characteristics along the way.

As often happens, this case started with a laudable goal. In 2007, Washington, D.C.'s new mayor, Adrian Fenty, was determined to turn around the city's underperforming schools. He had his work cut out for him: at the time, barely one out of every two high school students was surviving to graduation after ninth grade, and only 8 percent of eighth graders were performing at grade level in math. Fenty hired an education reformer named Michelle Rhee to fill a powerful new post, chancellor of Washington's schools.

The going theory was that the students weren't learning enough because their teachers weren't doing a good job. So in 2009, Rhee implemented a plan to weed out the low-performing teachers. This is the trend in troubled school districts around the country, and from a systems engineering perspective the thinking makes perfect sense: Evaluate the teachers. Get rid of the worst ones, and place the best ones where they can do the most good. In the language of data scientists, this "optimizes" the school system, presumably ensuring better results for the kids. Except for "bad" teachers, who could argue with that? Rhee developed a teacher assessment tool called IMPACT, and at the end of the 2009–10 school year the district fired all the teachers whose scores put them in the bottom 2 percent. At the end of the following year, another 5 percent, or 206 teachers, were booted out.

Sarah Wysocki, a fifth-grade teacher, didn't seem to have any reason to worry. She had been at MacFarland Middle School for only two years but was already getting excellent reviews from her principal and her students' parents. One evaluation praised her attentiveness to the children; another called her "one of the best teachers I've ever come into contact with."

Yet at the end of the 2010–11 school year, Wysocki received a miserable score on her IMPACT evaluation. Her problem was a new scoring system known as value-added modeling, which purported to measure her effectiveness in teaching math and language skills. That score, generated by an algorithm, represented half of her overall evaluation, and it outweighed the positive reviews from school administrators and the community. This left the district with no choice but to fire her, along with 205 other teachers who had IMPACT scores below the minimal threshold.

This didn't seem to be a witch hunt or a settling of scores. Indeed, there's a logic to the school district's approach. Administrators, after all, could be friends with terrible teachers. They could admire their style or their apparent dedication. Bad teachers can *seem* good. So Washington, like many other school systems, would minimize this human bias and pay more attention to scores based on hard results: achievement scores in math and reading. The numbers would speak clearly, district officials promised. They would be more fair.

Wysocki, of course, felt the numbers were horribly unfair, and she wanted to know where they came from. "I don't think anyone understood them," she later told me. How could a good teacher get such dismal scores? What was the value-added model measuring?

Well, she learned, it was complicated. The district had hired a consultancy, Princeton-based Mathematica Policy Research, to come up with the evaluation system. Mathematica's challenge was to measure the educational progress of the students in the district and then to calculate how much of their advance or decline could be attributed to their teachers. This wasn't easy, of course. The researchers knew that many variables, from students' socioeconomic backgrounds to the effects of learning disabilities, could affect student outcomes. The algorithms had to make allowances for such differences, which was one reason they were so complex.

Indeed, attempting to reduce human behavior, performance, and potential to algorithms is no easy job. To understand what Mathematica was up against, picture a ten-year-old girl living in a poor neighborhood in southeastern Washington, D.C. At the end of one school year, she takes her fifth-grade standardized test. Then life goes on. She may have family issues or money problems. Maybe she's moving from one house to another or worried about an older brother who's in trouble with the law. Maybe she's unhappy about her weight or frightened by a bully at school. In

any case, the following year she takes another standardized test, this one designed for sixth graders.

If you compare the results of the tests, the scores should stay stable, or hopefully, jump up. But if her results sink, it's easy to calculate the gap between her performance and that of the successful students.

But how much of that gap is due to her teacher? It's hard to know, and Mathematica's models have only a few numbers to compare. At Big Data companies like Google, by contrast, researchers run constant tests and monitor thousands of variables. They can change the font on a single advertisement from blue to red, serve each version to ten million people, and keep track of which one gets more clicks. They use this feedback to hone their algorithms and fine-tune their operation. While I have plenty of issues with Google, which we'll get to, this type of testing is an effective use of statistics.

Attempting to calculate the impact that one person may have on another over the course of a school year is much more complex. "There are so many factors that go into learning and teaching that it would be very difficult to measure them all," Wysocki says. What's more, attempting to score a teacher's effectiveness by analyzing the test results of only twenty-five or thirty students is statistically unsound, even laughable. The numbers are far too small given all the things that could go wrong. Indeed, if we were to analyze teachers with the statistical rigor of a search engine, we'd have to test them on thousands or even millions of randomly selected students. Statisticians count on large numbers to balance out exceptions and anomalies. (And WMDs, as we'll see, often punish individuals who happen to *be* the exception.)

Equally important, statistical systems require feedback—something to tell them when they're off track. Statisticians use errors to train their models and make them smarter. If Amazon.com,

through a faulty correlation, started recommending lawn care books to teenage girls, the clicks would plummet, and the algorithm would be tweaked until it got it right. Without feedback, however, a statistical engine can continue spinning out faulty and damaging analysis while never learning from its mistakes.

Many of the WMDs I'll be discussing in this book, including the Washington school district's value-added model, behave like that. They define their own reality and use it to justify their results. This type of model is self-perpetuating, highly destructive—and very common.

When Mathematica's scoring system tags Sarah Wysocki and 205 other teachers as failures, the district fires them. But how does it ever learn if it was right? It doesn't. The system itself has determined that they were failures, and that is how they are viewed. Two hundred and six "bad" teachers are gone. That fact alone appears to demonstrate how effective the value-added model is. It is cleansing the district of underperforming teachers. Instead of searching for the truth, the score comes to embody it.

This is one example of a WMD feedback loop. We'll see many of them throughout this book. Employers, for example, are increasingly using credit scores to evaluate potential hires. Those who pay their bills promptly, the thinking goes, are more likely to show up to work on time and follow the rules. In fact, there are plenty of responsible people and good workers who suffer misfortune and see their credit scores fall. But the belief that bad credit correlates with bad job performance leaves those with low scores less likely to find work. Joblessness pushes them toward poverty, which further worsens their scores, making it even harder for them to land a job. It's a downward spiral. And employers never learn how many good employees they've missed out on by focusing on credit scores. In WMDs, many poisonous assumptions are camouflaged by math and go largely untested and unquestioned.

This underscores another common feature of WMDs. They tend to punish the poor. This is, in part, because they are engineered to evaluate large numbers of people. They specialize in bulk, and they're cheap. That's part of their appeal. The wealthy, by contrast, often benefit from personal input. A white-shoe law firm or an exclusive prep school will lean far more on recommendations and face-to-face interviews than will a fast-food chain or a cash-strapped urban school district. The privileged, we'll see time and again, are processed more by people, the masses by machines.

Wysocki's inability to find someone who could explain her appalling score, too, is telling. Verdicts from WMDs land like dictates from the algorithmic gods. The model itself is a black box, its contents a fiercely guarded corporate secret. This allows consultants like Mathematica to charge more, but it serves another purpose as well: if the people being evaluated are kept in the dark, the thinking goes, they'll be less likely to attempt to game the system. Instead, they'll simply have to work hard, follow the rules, and pray that the model registers and appreciates their efforts. But if the details are hidden, it's also harder to question the score or to protest against it.

For years, Washington teachers complained about the arbitrary scores and clamored for details on what went into them. It's an algorithm, they were told. It's very complex. This discouraged many from pressing further. Many people, unfortunately, are intimidated by math. But a math teacher named Sarah Bax continued to push the district administrator, a former colleague named Jason Kamras, for details. After a back-and-forth that extended for months, Kamras told her to wait for an upcoming technical report. Bax responded: "How do you justify evaluating people by a measure for which you are unable to provide explanation?" But that's the nature of WMDs. The analysis is outsourced to

coders and statisticians. And as a rule, they let the machines do the talking.

Even so, Sarah Wysocki was well aware that her students' standardized test scores counted heavily in the formula. And here she had some suspicions. Before starting what would be her final year at MacFarland Middle School, she had been pleased to see that her incoming fifth graders had scored surprisingly well on their year-end tests. At Barnard Elementary School, where many of Sarah's students came from, 29 percent of the students were ranked at an "advanced reading level." This was five times the average in the school district.

Yet when classes started she saw that many of her students struggled to read even simple sentences. Much later, investigations by the *Washington Post* and *USA Today* revealed a high level of erasures on the standardized tests at forty-one schools in the district, including Barnard. A high rate of corrected answers points to a greater likelihood of cheating. In some of the schools, as many as 70 percent of the classrooms were suspected.

What does this have to do with WMDs? A couple of things. First, teacher evaluation algorithms are a powerful tool for behavioral modification. That's their purpose, and in the Washington schools they featured both a stick and a carrot. Teachers knew that if their students stumbled on the test their own jobs were at risk. This gave teachers a strong motivation to ensure their students passed, especially as the Great Recession battered the labor market. At the same time, if their students outperformed their peers, teachers and administrators could receive bonuses of up to $8,000. If you add those powerful incentives to the evidence in the case—the high number of erasures and the abnormally high test scores—there were grounds for suspicion that fourth-grade teachers, bowing either to fear or to greed, had corrected their students' exams.

It is conceivable, then, that Sarah Wysocki's fifth-grade students started the school year with artificially inflated scores. If so, their results the following year would make it appear that they'd lost ground in fifth grade—and that their teacher was an underperformer. Wysocki was convinced that this was what had happened to her. That explanation would fit with the observations from parents, colleagues, and her principal that she was indeed a good teacher. It would clear up the confusion. Sarah Wysocki had a strong case to make.

But you cannot appeal to a WMD. That's part of their fearsome power. They do not listen. Nor do they bend. They're deaf not only to charm, threats, and cajoling but also to logic—even when there is good reason to question the data that feeds their conclusions. Yes, if it becomes clear that automated systems are screwing up on an embarrassing and systematic basis, programmers will go back in and tweak the algorithms. But for the most part, the programs deliver unflinching verdicts, and the human beings employing them can only shrug, as if to say, "Hey, what can you do?"

And that is precisely the response Sarah Wysocki finally got from the school district. Jason Kamras later told the *Washington Post* that the erasures were "suggestive" and that the numbers might have been wrong in her fifth-grade class. But the evidence was not conclusive. He said she had been treated fairly.

Do you see the paradox? An algorithm processes a slew of statistics and comes up with a probability that a certain person *might* be a bad hire, a risky borrower, a terrorist, or a miserable teacher. That probability is distilled into a score, which can turn someone's life upside down. And yet when the person fights back, "suggestive" countervailing evidence simply won't cut it. The case must be ironclad. The human victims of WMDs, we'll see time and again, are held to a far higher standard of evidence than the algorithms themselves.

After the shock of her firing, Sarah Wysocki was out of a job for only a few days. She had plenty of people, including her principal, to vouch for her as a teacher, and she promptly landed a position at a school in an affluent district in northern Virginia. So thanks to a highly questionable model, a poor school lost a good teacher, and a rich school, which didn't fire people on the basis of their students' scores, gained one.

...

Following the housing crash, I woke up to the proliferation of WMDs in banking and to the danger they posed to our economy. In early 2011 I quit my job at the hedge fund. Later, after rebranding myself as a data scientist, I joined an e-commerce start-up. From that vantage point, I could see that legions of other WMDs were churning away in every conceivable industry, many of them exacerbating inequality and punishing the poor. They were at the heart of the raging data economy.

To spread the word about WMDs, I launched a blog, MathBabe. My goal was to mobilize fellow mathematicians against the use of sloppy statistics and biased models that created their own toxic feedback loops. Data specialists, in particular, were drawn to the blog, and they alerted me to the spread of WMDs in new domains. But in mid-2011, when Occupy Wall Street sprang to life in Lower Manhattan, I saw that we had work to do among the broader public. Thousands had gathered to demand economic justice and accountability. And yet when I heard interviews with the Occupiers, they often seemed ignorant of basic issues related to finance. They clearly hadn't been reading my blog. (I should add, though, that you don't need to understand all the details of a system to know that it has failed.)

I could either criticize them or join them, I realized, so I joined them. Soon I was facilitating weekly meetings of the Alternative

Banking Group at Columbia University, where we discussed financial reform. Through this process, I came to see that my two ventures outside academia, one in finance, the other in data science, had provided me with fabulous access to the technology and culture powering WMDs.

Ill-conceived mathematical models now micromanage the economy, from advertising to prisons. These WMDs have many of the same characteristics as the value-added model that derailed Sarah Wysocki's career in Washington's public schools. They're opaque, unquestioned, and unaccountable, and they operate at a scale to sort, target, or "optimize" millions of people. By confusing their findings with on-the-ground reality, most of them create pernicious WMD feedback loops.

But there's one important distinction between a school district's value-added model and, say, a WMD that scouts out prospects for extortionate payday loans. They have different payoffs. For the school district, the payoff is a kind of political currency, a sense that problems are being fixed. But for businesses it's just the standard currency: money. For many of the businesses running these rogue algorithms, the money pouring in seems to prove that their models are working. Look at it through their eyes and it makes sense. When they're building statistical systems to find customers or manipulate desperate borrowers, growing revenue appears to show that they're on the right track. The software is doing its job. The trouble is that profits end up serving as a stand-in, or proxy, for truth. We'll see this dangerous confusion crop up again and again.

This happens because data scientists all too often lose sight of the folks on the receiving end of the transaction. They certainly understand that a data-crunching program is bound to misinterpret people a certain percentage of the time, putting them in the wrong groups and denying them a job or a chance at their dream

house. But as a rule, the people running the WMDs don't dwell on those errors. Their feedback is money, which is also their incentive. Their systems are engineered to gobble up more data and fine-tune their analytics so that more money will pour in. Investors, of course, feast on these returns and shower WMD companies with more money.

And the victims? Well, an internal data scientist might say, no statistical system can be *perfect*. Those folks are collateral damage. And often, like Sarah Wysocki, they are deemed unworthy and expendable. Forget about them for a minute, they might say, and focus on all the people who get helpful suggestions from recommendation engines or who find music they love on Pandora, the ideal job on LinkedIn, or perhaps the love of their life on Match.com. Think of the astounding scale, and ignore the imperfections.

Big Data has plenty of evangelists, but I'm not one of them. This book will focus sharply in the other direction, on the damage inflicted by WMDs and the injustice they perpetuate. We will explore harmful examples that affect people at critical life moments: going to college, borrowing money, getting sentenced to prison, or finding and holding a job. All of these life domains are increasingly controlled by secret models wielding arbitrary punishments.

Welcome to the dark side of Big Data.

individual successes, often don't lead to lower health care spending. A 2013 study headed by Jill Horwitz, a law professor at UCLA, rips away the movement's economic underpinning. Randomized studies, according to the report, "raise doubts" that smokers and obese workers chalk up higher medical bills than others. While it is true that they are more likely to suffer from health problems, these tend to come later in life, when they're off the corporate health plan and on Medicare. In fact, the greatest savings from wellness programs come from the penalties assessed on the workers. In other words, like scheduling algorithms, they provide corporations with yet another tool to raid their employees' paychecks.

Despite my problems with wellness programs, they don't (yet) rank as full WMDs. They're certainly widespread, they intrude on the lives of millions of employees, and they inflict economic pain. But they are not opaque, and, except for the specious BMI score, they're not based on mathematical algorithms. They're a simple and widespread case of wage theft, one wrapped up in flowery health rhetoric.

Employers are already overdosing on our data. They're busy using it, as we've seen, to score us as potential employees and as workers. They're trying to map our thoughts and our friendships and predict our productivity. Since they're already deeply involved in insurance, with workforce health care a major expense, it's only natural that they would extend surveillance on a large scale to workers' health. And if companies cooked up their own health and productivity models, this could grow into a full-fledged WMD.

10

THE TARGETED CITIZEN

Civic Life

As you know by now, I am outraged by all sorts of WMDs. So let's imagine that I decide to launch a campaign for tougher regulations on them, and I post a petition on my Facebook page. Which of my friends will see it on their news feed?

I have no idea. As soon as I hit send, that petition belongs to Facebook, and the social network's algorithm makes a judgment about how to best use it. It calculates the odds that it will appeal to each of my friends. Some of them, it knows, often sign petitions, and perhaps share them with their own networks. Others tend to scroll right past. At the same time, a number of my friends pay more attention to me and tend to click the articles I post. The

Facebook algorithm takes all of this into account as it decides who will see my petition. For many of my friends, it will be buried so low on their news feed that they'll never see it.

This is what happens when the immensely powerful network we share with 1.5 billion users is also a publicly traded corporation. While Facebook may feel like a modern town square, the company determines, according to its own interests, what we see and learn on its social network. As I write this, about two-thirds of American adults have a profile on Facebook. They spend thirty-nine minutes a day on the site, only four minutes less than they dedicate to face-to-face socializing. Nearly half of them, according to a Pew Research Center report, count on Facebook to deliver at least some of their news, which leads to the question: By tweaking its algorithm and molding the news we see, can Facebook game the political system?

The company's own researchers have been looking into this. During the 2010 and 2012 elections, Facebook conducted experiments to hone a tool they called the "voter megaphone." The idea was to encourage people to spread word that they had voted. This seemed reasonable enough. By sprinkling people's news feeds with "I voted" updates, Facebook was encouraging Americans—more than sixty-one million of them—to carry out their civic duty and make their voices heard. What's more, by posting about people's voting behavior, the site was stoking peer pressure to vote. Studies have shown that the quiet satisfaction of carrying out a civic duty is less likely to move people than the possible judgment of friends and neighbors.

At the same time, Facebook researchers were studying how different types of updates influenced people's voting behavior. No researcher had ever worked in a human laboratory of this scale. Within hours, Facebook could harvest information from tens of millions of people, or more, measuring the impact that their

words and shared links had on each other. And it could use that knowledge to influence people's actions, which in this case happened to be voting.

That's a significant amount of power. And Facebook is not the only company to wield it. Other publicly held corporations, including Google, Apple, Microsoft, Amazon, and cell phone providers like Verizon and AT&T, have vast information on much of humanity—and the means to steer us in any way they choose.

Usually, as we've seen, they're focused on making money. However, their profits are tightly linked to government policies. The government regulates them, or chooses not to, approves or blocks their mergers and acquisitions, and sets their tax policies (often turning a blind eye to the billions parked in offshore tax havens). This is why tech companies, like the rest of corporate America, inundate Washington with lobbyists and quietly pour hundreds of millions of dollars in contributions into the political system. Now they're gaining the wherewithal to fine-tune our political behavior—and with it the shape of American government—just by tweaking their algorithms.

The Facebook campaign started out with a constructive and seemingly innocent goal: to encourage people to vote. And it succeeded. After comparing voting records, researchers estimated that their campaign had increased turnout by 340,000 people. That's a big enough crowd to swing entire states, and even national elections. George W. Bush, after all, won in 2000 by a margin of 537 votes in Florida. The activity of a single Facebook algorithm on Election Day, it's clear, could not only change the balance of Congress but also decide the presidency.

Facebook's potency comes not only from its reach but also from its ability to use its own customers to influence their friends. The vast majority of the sixty-one million people in the experiment received a message on their news feed encouraging them to vote.

The message included a display of photos: six of the user's Facebook friends, randomly selected, who had clicked the "I Voted" button. The researchers also studied two control groups, each numbering around six hundred thousand. One group saw the "I Voted" campaign, but without the pictures of friends. The other received nothing at all.

By sprinkling its messages through the network, Facebook was studying the impact of friends' behavior on our own. Would people encourage their friends to vote, and would this affect their behavior? According to the researchers' calculations, seeing that friends were participating made all the difference. People paid much more attention when the "I Voted" updates came from friends, and they were more likely to share those updates. About 20 percent of the people who saw that their friends had voted also clicked on the "I Voted" button. Among those who didn't get the button from friends, only 18 percent did. We can't be sure that all the people who clicked the button actually voted, or that those who didn't click it stayed home. Still, with sixty-one million potential voters on the network, a possible difference of two points can be huge.

Two years later, Facebook took a step further. For three months leading up to the election between President Obama and Mitt Romney, a researcher at the company, Solomon Messing, altered the news feed algorithm for about two million people, all of them politically engaged. These people got a higher proportion of hard news, as opposed to the usual cat videos, graduation announcements, or photos from Disney World. If their friends shared a news story, it showed up high on their feed.

Messing wanted to see if getting more news from friends changed people's political behavior. Following the election, Messing sent out surveys. The self-reported results indicated that the voter participation in this group inched up from 64 to 67 percent.

"When your friends deliver the newspaper," said Lada Adamic, a computational social scientist at Facebook, "interesting things happen." Of course, it wasn't really the friends delivering the newspaper, but Facebook itself. You might argue that newspapers have exerted similar power for eons. Editors pick the front-page news and decide how to characterize it. They choose whether to feature bombed Palestinians or mourning Israelis, a policeman rescuing a baby or battering a protester. These choices can no doubt influence both public opinion and elections. The same goes for television news. But when the *New York Times* or CNN covers a story, everyone sees it. Their editorial decision is clear, on the record. It is not opaque. And people later debate (often on Facebook) whether that decision was the right one.

Facebook is more like the Wizard of Oz: we do not see the human beings involved. When we visit the site, we scroll through updates from our friends. The machine appears to be only a neutral go-between. Many people still believe it is. In 2013, when a University of Illinois researcher named Karrie Karahalios carried out a survey on Facebook's algorithm, she found that 62 percent of the people were unaware that the company tinkered with the news feed. They believed that the system instantly shared everything they posted with all of their friends.

The potential for Facebook to hold sway over our politics extends beyond its placement of news and its Get Out the Vote campaigns. In 2012, researchers experimented on 680,000 Facebook users to see if the updates in their news feeds could affect their mood. It was already clear from laboratory experiments that moods are contagious. Being around a grump is likely to turn you into one, if only briefly. But would such contagions spread online?

Using linguistic software, Facebook sorted positive (stoked!) and negative (bummed!) updates. They then reduced the volume of downbeat postings in half of the news feeds, while reducing the

cheerful quotient in the others. When they studied the users' subsequent posting behavior, they found evidence that the doctored new feeds had indeed altered their moods. Those who had seen fewer cheerful updates produced more negative posts. A similar pattern emerged on the positive side.

Their conclusion: "Emotional states can be transferred to others . . . , leading people to experience the same emotions without their awareness." In other words, Facebook's algorithms can affect how millions of people feel, and those people won't know that it's happening. What would occur if they played with people's emotions on Election Day?

I have no reason to believe that the social scientists at Facebook are actively gaming the political system. Most of them are serious academics carrying out research on a platform that they could only have dreamed about two decades ago. But what they have demonstrated is Facebook's enormous power to affect what we learn, how we feel, and whether we vote. Its platform is massive, powerful, and opaque. The algorithms are hidden from us, and we see only the results of the experiments researchers choose to publish.

Much the same is true of Google. Its search algorithm appears to be focused on raising revenue. But search results, if Google so chose, could have a dramatic effect on what people learn and how they vote. Two researchers, Robert Epstein and Ronald E. Robertson, recently asked undecided voters in both the United States and India to use a search engine to learn about upcoming elections. The engines they used were programmed to skew the search results, favoring one party over another. Those results, they said, shifted voting preferences by 20 percent.

This effect was powerful, in part, because people widely trust search engines. Some 73 percent of Americans, according to a Pew Research report, believe that search results are both accurate and impartial. So companies like Google would be risking their

own reputation, and inviting a regulatory crackdown, if they doctored results to favor one political outcome over another.

Then again, how would anyone know? What we learn about these Internet giants comes mostly from the tiny proportion of their research that they share. Their algorithms represent vital trade secrets. They carry out their business in the dark.

I wouldn't yet call Facebook or Google's algorithms political WMDs, because I have no evidence that the companies are using their networks to cause harm. Still, the potential for abuse is vast. The drama occurs in code and behind imposing firewalls. And as we'll see, these technologies can place each of us into our own cozy political nook.

• • •

By late spring of 2012, the former governor of Massachusetts, Mitt Romney, had the Republican nomination sewn up. The next step was to build up his war chest for the general election showdown with President Obama. And so on May 17, he traveled to Boca Raton, Florida, for a fund-raiser at the palatial home of Marc Leder, a private equity investor. Leder had already poured $225,000 into the pro-Romney Super PAC Restore Our Future and had given another $63,330 to the Romney Victory PAC. He had gathered a host of rich friends, most of them in finance and real estate, to meet the candidate. Naturally, the affair would be catered.

Romney could safely assume that he was walking into a closed setting with a group of people who thought much like Marc Leder. If this had been a televised speech, Romney would have taken great care not to ruffle potential Republican voters. Those ranged from Evangelical Christians and Wall Street financiers to Cuban Americans and suburban soccer moms. Trying to please everyone is one reason most political speeches are boring (and Romney's, even his supporters groused, were especially so). But at

an intimate gathering at Marc Leder's house, a small and influential group might get closer to the real Mitt Romney and hear what the candidate really believed, unfiltered. They had already given him large donations. A frank chat was the least they could expect for their investment.

Basking in the company of people he believed to be supportive and like-minded, Romney let loose with his observation that 47 percent of the population were "takers," living off the largesse of big government. These people would never vote for him, the governor said—which made it especially important to reach out to the other 53 percent. But Romney's targeting, it turned out, was inexact. The caterers circulating among the donors, serving drinks and canapés, were outsiders. And like nearly everyone in the developed world, they carried phones equipped with video cameras. Romney's dismissive remarks, captured by a bartender, went viral. The gaffe very likely cost Romney any chance he had of winning the White House.

Success for Romney at that Boca Raton gathering required both accurate targeting and secrecy. He wanted to be the ideal candidate for Marc Leder and friends. And he trusted that Leder's house represented a safe zone in which to be that candidate. In a dream world, politicians would navigate countless such targeted safe zones so that they could tailor their pitch for every subgroup—without letting the others see it. One candidate could be many candidates, with each part of the electorate seeing only the parts they liked.

This duplicity, or "multiplicity," is nothing new in politics. Politicians have long tried to be many things to many people, whether they're eating kielbasa in Milwaukee, quoting the Torah in Brooklyn, or pledging allegiance to corn-based ethanol in Iowa. But as Romney discovered, video cameras can now bust them if they overdo their contortions.

Modern consumer marketing, however, provides politicians with new pathways to specific voters so that they can tell them what they know they want to hear. Once they do, those voters are likely to accept the information at face value because it confirms their previous beliefs, a phenomenon psychologists call confirmation bias. It is one reason that none of the invited donors at the Romney event questioned his assertion that nearly half of voters were hungry for government handouts. It only bolstered their existing beliefs.

This merging of politics and consumer marketing has been developing for the last half century, as the tribal rituals of American politics, with their ward bosses and long phone lists, have given way to marketing science. In *The Selling of the President*, which followed Richard Nixon's 1968 campaign, the journalist Joe McGinniss introduced readers to the political operatives working to market the presidential candidate like a consumer good. By using focus groups, Nixon's campaign was able to hone his pitch for different regions and demographics.

But as time went on, politicians wanted a more detailed approach, one that would ideally reach each voter with a personalized come-on. This desire gave birth to direct-mail campaigns. Borrowing tactics from the credit card industry, political operatives built up huge databases of customers—voters, in this case—and placed them into various subgroups, reflecting their values and their demographics. For the first time, it was possible for next-door neighbors to receive different letters or brochures from the same politician, one vowing to protect wilderness and the other stressing law and order.

Direct mail was microtargeting on training wheels. The convergence of Big Data and consumer marketing now provides politicians with far more powerful tools. They can target micro-groups of citizens for both votes and money and appeal to each of

them with a meticulously honed message, one that no one else is likely to see. It might be a banner on Facebook or a fund-raising email. But each one allows candidates to quietly sell multiple versions of themselves—and it's anyone's guess which version will show up for work after inauguration.

• • •

In July of 2011, more than a year before President Obama would run for reelection, a data scientist named Rayid Ghani posted an update on LinkedIn:

> *Hiring analytics experts who want to make a difference. The Obama re-election campaign is growing the analytics team to work on high-impact large-scale data mining problems.*
>
> *We have several positions available at all levels of experience. Looking for experts in statistics, machine learning, data mining, text analytics, and predictive analytics to work with large amounts of data and help guide election strategy.*

Ghani, a computer scientist educated at Carnegie Mellon, would be heading up the data team for Obama's campaign. In his previous position, at Accenture Labs in Chicago, Ghani had developed consumer applications for Big Data, and he trusted that he could apply his skills to politics. The goal for the Obama campaign was to create tribes of like-minded voters, people as uniform in their values and priorities as the guests at Marc Leder's reception—but without the caterers. Then they could target them with the messaging most likely to move them toward specific objectives, including voting, organizing, and fund-raising.

One of Ghani's projects at Accenture involved modeling supermarket shoppers. A major grocer had provided the Accenture team with a massive database of anonymized consumer purchases. The idea was to dig into this data to study each consumer's buying habits and then to place the shoppers into hundreds of different consumer buckets. There would be the impulse shoppers who bought candy at the checkout counter and the health nuts who were willing to pay triple for organic kale. Those were the obvious categories. But others were more surprising. Ghani and his team, for example, could spot people who stuck close to a brand and others who would switch for even a tiny discount. There were buckets for these "persuadables," too. The end goal was to come up with a different plan for each shopper and to guide them through the store, leading them to all the foods they were most likely to want and buy.

Unfortunately for Accenture's clients, this ultimate vision hinged upon the advent of computerized shopping carts, which haven't yet caught on in a big way and maybe never will. But despite the disappointment in supermarkets, Ghani's science translated perfectly into politics. Those fickle shoppers who switched brands to save a few cents, for example, behaved very much like swing voters. In the supermarket, it was possible to estimate how much it would cost to turn each shopper from one brand of ketchup or coffee to another more profitable brand. The supermarket could then pick out, say, the 15 percent most likely to switch and provide them with coupons. Smart targeting was essential. They certainly didn't want to give coupons to shoppers who were ready to pay full price. That was like burning money.*

Would similar calculations work for swing voters? Armed with massive troves of consumer, demographic, and voting data, Ghani

* Similarly, consumer websites are much more likely to offer discounts to people who are not already logged in. This is another reason to clear your cookies regularly.

and his team set out to investigate. However, they faced one crucial difference. In the supermarket project, all of the available data related precisely to the shopping domain. They studied shopping patterns to predict (and influence) what people would buy. But in politics there was very little relevant data available. Data teams for both campaigns needed proxies, and this required research.

They started out by interviewing several thousand people in great depth. These folks fell into different groups. Some cared about education or gay rights, others worried about Social Security or the impact of fracking on freshwater aquifers. Some supported the president unconditionally. Others sat on the fence. A good number liked him but didn't usually get around to voting. Some of them—and this was vital—were ready to contribute money to Obama's campaign.

Once Ghani's data team understood this small group of voters, their desires, their fears, and what it took to change their behavior, the next challenge was to find millions of other voters (and donors) who resembled them. This involved plowing through the consumer data and demographics of the voters they had interviewed and building mathematical profiles of them. Then it was just a matter of scouring national databases, finding people with similar profiles, and placing them into the same buckets.

The campaign could then target each group with advertisements, perhaps on Facebook or the media sites they visited, to see if they responded as expected. They carried out the same kind of A/B testing that Google uses to see which shade of blue garners more clicks on a button. Trying different approaches, they found, for example, that e-mail subject lines reading only "Hey!" bugged people but also led to more engagement and sometimes more donations. Through thousands of tests and tweaks, the campaign finally sized up its audience—including an all-important contingent of fifteen million swing voters.

Throughout this process, each campaign developed profiles of American voters. Each profile contained numerous scores, which not only gauged their value as a potential voter, volunteer, and donor but also reflected their stances on different issues. One voter might have a high score on environmental issues but a low one on national security or international trade. These political profiles are very similar to those that Internet companies, like Amazon and Netflix, use to manage their tens of millions of customers. Those companies' analytics engines churn out nearly constant cost/benefit analyses to maximize their revenue per customer.

Four years later, Hillary Clinton's campaign built upon the methodology established by Obama's team. It contracted a micro-targeting start-up, the Groundwork, financed by Google chairman Eric Schmidt and run by Michael Slaby, the chief technology officer of Obama's 2012 campaign. The goal, according to a report in Quartz, was to build a data system that would create a political version of systems that companies like Salesforce.com develop to manage their millions of customers.

The appetite for fresh and relevant data, as you might imagine, is intense. And some of the methods used to gather it are unsavory, not to mention intrusive. In late 2015, the *Guardian* reported that a political data firm, Cambridge Analytica, had paid academics in the United Kingdom to amass Facebook profiles of US voters, with demographic details and records of each user's "likes." They used this information to develop psychographic analyses of more than forty million voters, ranking each on the scale of the "big five" personality traits: openness, conscientiousness, extroversion, agreeableness, and neuroticism. Groups working with the Ted Cruz presidential campaign then used these studies to develop television commercials targeted for different types of voters, placing them in programming they'd be most likely to watch. When the Republican Jewish Coalition was meeting at the Venetian

in Las Vegas in May 2015, for instance, the Cruz campaign unleashed a series of web-based advertisements visible only inside the hotel complex that emphasized Cruz's devotion to Israel and its security.

I should mention here that not all of these targeting campaigns have proven to be effective. Some, no doubt, are selling little more than snake oil. The microtargeters, after all, are themselves marketing to campaigns and political action groups with millions of dollars to spend. They sell them grand promises of priceless databases and pinpoint targeting, many of which are bound to be exaggerated. So in this sense the politicians not only purvey questionable promises but also consume them (at exorbitant expense). That said, as the Obama team demonstrated, some of these methods are fruitful. And so the industry—serious data scientists and hucksters alike—zeros in on voters.

Political microtargeters, however, face unique constraints, which make their work far more complex. The value of each voter, for example, rises or falls depending on the probability that his or her state will be in play. A swing voter in a swing state, like Florida, Ohio, or Nevada, is highly valuable. But if polls show the state tilting decisively to either blue or red, that voter's value plummets, and the marketing budget is quickly shifted toward other voters whose value is climbing.

In this sense, we can think of the voting public very much as we think of financial markets. With the flow of information, values rise and fall, as do investments. In these new political markets, each one of us represents a stock with its own fluctuating price. And each campaign must decide if and how to invest in us. If we merit the investment, then they decide not only what information to feed us but also how much and how to deliver it.

Similar calculations, on a macro scale, have been going on for decades, as campaigns plot their TV spending. As polling numbers change, they might cut ads in Pittsburgh and move those dollars to Tampa or Las Vegas. But with microtargeting, the focus shifts from the region to the individual. More important, that individual alone sees the customized version of the politician.

The campaigns use similar analysis to identify potential donors and to optimize each one. Here it gets complicated, because many of the donors themselves are carrying out their own calculations. They want the biggest bang for their buck. They know that if they immediately hand over the maximum contribution the campaign will view them as "fully tapped" and therefore irrelevant. But refusing to give any money will also render them irrelevant. So many give a drip-feed of money based on whether the messages they hear are ones they agree with. For them, managing a politician is like training a dog with treats. This training effect is all the more powerful for contributors to Super PACS, which do not limit political contributions.

The campaigns, of course, are well aware of this tactic. With microtargeting, they can send each of those donors the information most likely to pry more dollars from their bank accounts. And these messages will vary from one donor to the next.

• • •

These tactics aren't limited to campaigns. They infect our civic life, with lobbyists and interest groups now using these targeting methods to carry out their dirty work. In 2015, the Center for Medical Progress, an antiabortion group, posted videos featuring what they claimed was an aborted fetus at a Planned Parenthood clinic. The videos asserted that Planned Parenthood doctors were selling baby parts for research, and they spurred a wave of protest, and a Republican push to eliminate the organization's funding.

Research later showed that the video had been doctored: the so-called fetus was actually a photo of a stillborn baby born to a

woman in rural Pennsylvania. And Planned Parenthood does not sell fetal tissue. The Center for Medical Progress admitted that the video contained misinformation. That weakened its appeal for a mass market. But with microtargeting, antiabortion activists could continue to build an audience for the video, despite the flawed premise, and use it to raise funds to fight Planned Parenthood.

While that campaign launched into public view, hundreds of others continue to hover below the surface, addressing individual voters. These quieter campaigns are equally deceptive and even less accountable. And they deliver ideological bombs that politicians will only hint at on the record. According to Zeynep Tufekci, a techno-sociologist and professor at the University of North Carolina, these groups pinpoint vulnerable voters and then target them with fear-mongering campaigns, scaring them about their children's safety or the rise of illegal immigration. At the same time, they can keep those ads from the eyes of voters likely to be turned off (or even disgusted) by such messaging.

Successful microtargeting, in part, explains why in 2015 more than 43 percent of Republicans, according to a survey, still believed the lie that President Obama is a Muslim. And 20 percent of Americans believed he was born outside the United States and, consequently, an illegitimate president. (Democrats may well spread their own disinformation in microtargeting, but nothing that has surfaced matches the scale of the anti-Obama campaigns.)

Even with the growth of microtargeting, political campaigns are still directing 75 percent of their media buy, on average, to television. You might think that this would have an equalizing effect, and it does. Television delivers the broader, and accountable, messaging, while microtargeting does its work in the shadows. But even television is moving toward personalized advertising. New advertising companies like Simulmedia, in New York, assemble

TV viewers into behavioral buckets, so that advertisers can target audiences of like-minded people, whether hunters, pacifists, or buyers of tank-sized SUVs. As television and the rest of the media move toward profiling their viewers, the potential for political microtargeting grows.

As this happens, it will become harder to access the political messages our neighbors are seeing—and as a result, to understand why they believe what they do, often passionately. Even a nosy journalist will struggle to track down the messaging. It is not enough simply to visit the candidate's web page, because they, too, automatically profile and target each visitor, weighing everything from their zip codes to the links they click on the page, even the photos they appear to look at. It's also fruitless to create dozens of "fake" profiles, because the systems associate each real voter with deep accumulated knowledge, including purchasing records, addresses, phone numbers, voting records, and even social security numbers and Facebook profiles. To convince the system it's real, each fake would have to come with its own load of data. Fabricating one would require far too much work for a research project (and in the worst-case scenario it might get the investigator tangled up in fraud).

The result of these subterranean campaigns is a dangerous imbalance. The political marketers maintain deep dossiers on us, feed us a trickle of information, and measure how we respond to it. But we're kept in the dark about what our neighbors are being fed. This resembles a common tactic used by business negotiators. They deal with different parties separately so that none of them knows what the other is hearing. This asymmetry of information prevents the various parties from joining forces—which is precisely the point of a democratic government.

This growing science of microtargeting, with its profiles and predictions, fits all too neatly into our dark collection of WMDs.

It is vast, opaque, and unaccountable. It provides cover to politicians, encouraging them to be many things to many people.

The scoring of individual voters also undermines democracy, making a minority of voters important and the rest little more than a supporting cast. Indeed, looking at the models used in presidential elections, we seem to inhabit a shrunken country. As I write this, the entire voting population that *matters* lives in a handful of counties in Florida, Ohio, Nevada, and a few other swing states. Within those counties is a small number of voters whose opinions weigh in the balance. I might point out here that while many of the WMDs we've been looking at, from predatory ads to policing models, deliver most of their punishment to the struggling classes, political microtargeting harms voters of every economic class. From Manhattan to San Francisco, rich and poor alike find themselves disenfranchised (though the truly affluent, of course, can more than compensate for this with campaign contributions).

In any case, the entire political system—the money, the attention, the fawning—turns to targeted voters like a flower following the sun. The rest of us are virtually ignored (except for fund-raising come-ons). The programs have already predicted our voting behavior, and any attempt to change it is not worth the investment.*

This creates a nefarious feedback loop. The disregarded voters are more likely to grow disenchanted. The winners know how to play the game. They get the inside story, while the vast majority of consumers receive only market-tested scraps.

* At the federal level, this problem could be greatly alleviated by abolishing the Electoral College system. It's the winner-take-all mathematics from state to state that delivers so much power to a relative handful of voters. It's as if in politics, as in economics, we have a privileged 1 percent. And the money from the financial 1 percent underwrites the microtargeting to secure the votes of the political 1 percent. Without the Electoral College, by contrast, every vote would be worth exactly the same. That would be a step toward democracy.

Indeed, there is an added asymmetry. People who are expected to be voters but who, for one reason or another, skip an election find themselves lavished with attention the next time round. They still seem to brim with high voting potential. But those expected not to vote are largely ignored. The systems are searching for the cheapest votes to convert, with the highest return for each dollar spent. And nonvoters often look expensive. This dynamic prods a certain class of people to stay active and lets the rest lie fallow forever.

As is often the case with WMDs, the very same models that inflict damage could be used to humanity's benefit. Instead of targeting people in order to manipulate them, it could line them up for help. In a mayoral race, for example, a microtargeting campaign might tag certain voters for angry messages about unaffordable rents. But if the candidate knows these voters are angry about rent, how about using the same technology to identify the ones who will most benefit from affordable housing and then help them find it?

With political messaging, as with most WMDs, the heart of the problem is almost always the objective. Change that objective from leeching off people to helping them, and a WMD is disarmed—and can even become a force for good.