




WORLD WITHOUT MIND

THE EXISTENTIAL THREAT OF BIG TECH

FRANKLIN FOER

PENGUIN PRESS
NEW YORK
2017

Three

MARK ZUCKERBERG'S WAR ON FREE WILL

SILICON VALLEY GRADUATED from the counterculture, but not really. All the values it professes are the values of the sixties. The big tech companies present themselves as platforms for personal liberation, just as Stewart Brand preached. Everyone has the right to speak their mind on social media, to fulfill their intellectual and democratic potential, to express their individuality. Where television had been a passive medium that rendered citizens inert, Facebook is participatory and empowering. It allows users to read widely, think for themselves, and form their own opinions.

We can't entirely dismiss this rhetoric. There are parts of the world, even in the United States, where Facebook emboldens citizens and enables them to organize themselves in opposition to power. But we shouldn't accept Facebook's self-conception as sincere, either. Facebook is a carefully managed top-down system, not a robust public square. It mimics some of the patterns of conversation, but that's a surface trait. In reality, Facebook is a tangle

MARK ZUCKERBERG'S WAR ON FREE WILL

of rules and procedures for sorting information, rules devised by the corporation for the ultimate benefit of the corporation. Facebook is always surveilling users, always auditing them, using them as lab rats in its behavioral experiments. While it creates the impression that it offers choice, Facebook paternalistically nudges users in the direction it deems best for them, which also happens to be the direction that thoroughly addicts them. It's a phoniness most obvious in the compressed, historic career of Facebook's mastermind.

↓ Start new

MARK ZUCKERBERG IS A GOOD BOY, but he wanted to be bad, or maybe just a little bit naughty. The heroes of his adolescence were the original hackers. Let's be precise about the term. His idols weren't malevolent data thieves or cyber-terrorists. In the parlance of hacker culture, such ill-willed outlaws are known as crackers. Zuckerberg never put crackers on a pedestal. Still, his hacker heroes were disrespectful of authority. They were technically virtuosic, infinitely resourceful nerd cowboys, unbound by conventional thinking. In MIT's labs, during the sixties and seventies, they broke any rule that interfered with building the stuff of early computing, such marvels as the first video games and word processors. With their free time, they played epic pranks, which happened to draw further attention to their own cleverness—installing a living, breathing cow on the roof of a Cambridge dorm; launching a weather balloon, which miraculously emerged from beneath the turf, emblazoned with "MIT," in the middle of a Harvard-Yale football game.

The hackers' archenemies were the bureaucrats who ran universities, corporations, and governments. Bureaucrats talked about

making the world more efficient, just like the hackers. But they were really small-minded paper-pushers who fiercely guarded the information they held, even when that information yearned to be shared. When hackers clearly engineered better ways of doing things—a box that enabled free long-distance calls, an instruction that might improve an operating system—the bureaucrats stood in their way, wagging an unbending finger. The hackers took aesthetic and comic pleasure in outwitting the men in suits.

When Zuckerberg arrived at Harvard in the fall of 2002, the heyday of the hackers had long passed. They were older guys now, the stuff of good tales, some stuck in twilight struggles against The Man. But Zuckerberg wanted to hack, too, and with that old-time indifference to norms. In high school—using the *nom de hack* Zuck Fader—he picked the lock that prevented outsiders from fiddling with AOL's code and added his own improvements to its instant messaging program. As a college sophomore he hatched a site called Facemash—with the high-minded purpose of determining the hottest kid on campus. Zuckerberg asked users to compare images of two students and then determine the better looking of the two. The winner of each pairing advanced to the next round of his hormonal tournament. To cobble this site together, Zuckerberg needed photos. He purloined those from the servers of the various Harvard houses that stockpiled them. "One thing is certain," he wrote on a blog as he put the finishing touches on his creation, "and it's that I'm a jerk for making this site. Oh well."

His brief experimentation with rebellion ended with his apologizing to a Harvard disciplinary panel, as well as campus women's groups, and mulling strategies to redeem his soiled reputation. In the years since, he's shown that defiance really wasn't his natu-

ral inclination. His distrust of authority was such that he sought out Don Graham, then the venerable chairman of the Washington Post company, as his mentor. After he started Facebook, he shadowed various giants of corporate America so that he could study their managerial styles up close. Though he hasn't fully shed his awkward ways, he has sufficiently overcome his introversion to appear at fancy dinner parties, Charlie Rose interviews, and *Vanity Fair* cover shoots.

Still, the juvenile fascination with hackers never did die, or rather he carried it forward into his new, more mature incarnation. When he finally had a corporate campus of his own, he procured a vanity address for it: One Hacker Way. He designed a plaza with h-a-c-k inlaid into the concrete. In the center of his office park, he created an open meeting space called Hacker Square. This is, of course, the venue where his employees join for all-night Hackathons. As he told a group of would-be entrepreneurs, "We've got this whole ethos that we want to build a hacker culture."

Plenty of companies have similarly appropriated hacker culture—hackers are the ur-disrupters—but none have gone as far as Facebook. Of course, that's not without risks. "Hacking" is a loaded term, and a potentially alienating one, at least to shareholders who crave sensible rule-abiding leadership. But by the time Zuckerberg began extolling the virtues of hacking, he'd stripped the name of most of its original meaning and distilled it into a managerial philosophy that contains barely a hint of rebelliousness. It might even be the opposite of rebelliousness. Hackers, he told one interviewer, were "just this group of computer scientists who were trying to quickly prototype and see what was possible. That's what I try to encourage our engineers to do here."

To hack is to be a good worker, a responsible Facebook citizen—a microcosm of the way in which the company has taken the language of radical individualism and deployed it in the service of conformism.

Zuckerberg claimed to have distilled that hacker spirit into a motivational motto: “Move Fast and Break Things.” Indeed, Facebook has excelled at that. The truth is, Facebook moved faster than Zuckerberg could ever have imagined. He hadn’t really intended his creation. His company was, as we all know, a dorm room lark, a thing he ginned up in a Red Bull-induced fit of sleeplessness. As his creation grew, it needed to justify its new scale to its investors, to its users, to the world. It needed to grow up fast. According to Dustin Moskovitz, who cofounded the company with Zuckerberg at Harvard, “It was always very important for our brand to get away from the image of frivolity it had, especially in Silicon Valley.” Over the span of its short life, the company has caromed from self-description to self-description. It has called itself a tool, a utility, and a platform. It has talked about openness and connectedness. And in all these attempts at defining itself, it has managed to clarify its intentions.

Though Facebook will occasionally talk about the transparency of governments and corporations, what it really wants to advance is the transparency of individuals—or what it has called, at various moments, “radical transparency” or “ultimate transparency.” The theory holds that the sunshine of sharing our intimate details will disinfect the moral mess of our lives. Even if we don’t intend for our secrets to become public knowledge, their exposure will improve society. With the looming threat that our embarrassing information will be broadcast, we’ll behave better. And perhaps the ubiquity of incriminating photos and damning reve-

lations will prod us to become more tolerant of one another’s sins. Besides, there’s virtue in living our lives truthfully. “The days of you having a different image for your work friends or co-workers and for the other people you know are probably coming to an end pretty quickly,” Zuckerberg has said. “Having two identities for yourself is an example of a lack of integrity.”

The point is that Facebook has a strong, paternalistic view on what’s best for you, and it’s trying to transport you there. “To get people to this point where there’s more openness—that’s a big challenge. But I think we’ll do it,” Zuckerberg has said. He has reason to believe that he will achieve that goal. With its size, Facebook has amassed outsized powers. These powers are so great that Zuckerberg doesn’t bother denying that fact. “In a lot of ways Facebook is more like a government than a traditional company. We have this large community of people, and more than other technology companies we’re really setting policies.”

WITHOUT KNOWING IT, Zuckerberg is the heir to a long political tradition. Over the last two hundred years, the West has been unable to shake an abiding fantasy, a dream sequence in which we throw out the bum politicians and replace them with engineers—rule by slide rule. The French were the first to entertain this notion in the bloody, world-churning aftermath of their revolution. A coterie of the country’s most influential philosophers (notably, Henri de Saint-Simon and Auguste Comte) were genuinely torn about the course of the country. They hated all the old ancient bastions of parasitic power—the feudal lords, the priests, and the warriors—but they also feared the chaos of the mob. To split the difference, they proposed a form of technocracy—engineers and

assorted technicians would rule with beneficent disinterestedness. Engineers would strip the old order of its power, while governing in the spirit of science. They would impose rationality and order.

This dream has captivated intellectuals ever since, especially Americans. The great sociologist Thorstein Veblen was obsessed with installing engineers in power and, in 1921, wrote a book making his case. His vision briefly became a reality. In the aftermath of World War I, American elites were aghast at all the irrational impulses unleashed by that conflict—the xenophobia, the racism, the urge to lynch and riot. What's more, the realities of economic life had grown so complicated, how could politicians possibly manage them? Americans of all persuasions began yearning for the salvific ascendance of the most famous engineer of his time: Herbert Hoover. During the war, Hoover had organized a system that managed to feed starving Europe, despite the seeming impossibility of that assignment. In 1920, Franklin Roosevelt—who would, of course, ultimately vanquish him from politics—organized a movement to draft Hoover for the presidency.

The Hoover experiment, in the end, hardly realized the happy fantasies about the Engineer King. A very different version of this dream, however, has come to fruition, in the form of the CEOs of the big tech companies. We're not ruled by engineers, not yet, but they have become the dominant force in American life, the highest, most influential tier of our elite. Marc Andreessen coined a famous aphorism that holds, "Software is eating the world." There's a bit of obfuscation in that formula—it's really the authors of software who are eating the world.

There's another way to describe this historical progression. Automation has come in waves. During the Industrial Revolu-

tion, machinery replaced manual workers. At first machines required human operators. Over time, machines came to function with hardly any human intervention. For centuries, engineers automated physical labor; our new engineering elite has automated thought. They have perfected technologies that take over intellectual processes, that render the brain redundant. Or as Marissa Mayer once argued, "You have to make *words* less human and more a piece of the *machine*." Indeed, we have begun to outsource our intellectual work to companies that suggest what we should learn, the topics we should consider, and the items we ought to buy. These companies can justify their incursions into our lives with the very arguments that Saint-Simon and Comte articulated: They are supplying us with efficiency; they are imposing order on human life.

Nobody better articulates the modern faith in engineering's power to transform society than Zuckerberg. He told a group of software developers, "You know, I'm an engineer, and I think a key part of the engineering mindset is this hope and this belief that you can take any system that's out there and make it much, much better than it is today. Anything, whether it's hardware, or software, a company, a developer ecosystem, you can take anything and make it much, much better." The world will improve, if only Zuckerberg's reason can prevail—and it will.

THE PRECISE SOURCE OF FACEBOOK'S power is algorithms. That's a concept repeated dutifully in nearly every story about the tech giants, yet it remains fuzzy at best to users of those sites. From the moment of the algorithm's invention, it was possible to see its power, its revolutionary potential. The algorithm was developed in order

to automate thinking, to remove difficult decisions from the hands of humans, to settle contentious debates. To understand the essence of the algorithm—and its utopian pretension—it's necessary to travel back to its birthplace, the brain of one of history's unimpeachable geniuses, Gottfried Leibniz.

Fifty years younger than Descartes, Leibniz grew up in the same world of religious conflict. His native Germany, Martin Luther's homeland, had become one of history's most horrific abattoirs, the contested territory at the center of the Thirty Years War. Although the battlefield made its own contribution to the corpse count, the aftermath of war was terrible, too. Dysentery, typhus, and plague conquered the German principalities. Famine and demographic collapse followed battle, some four million deaths in total. The worst-clobbered of the German states lost more than half of their population.

Leibniz was born as Europe negotiated the Peace of Westphalia ending the slaughter, so it was inevitable that he trained his prodigious intellectual energies on reconciling Protestants and Catholics, crafting schemes to unify humanity. Prodigious is perhaps an inadequate term to describe Leibniz's mental reserves. He produced schemes at, more or less, the rate he contracted his diaphragm. His archives, which still haven't been fully published, contain some two hundred thousand pages of his writing, filled with spectacular creations. Leibniz invented calculus—to be sure, he hadn't realized that Newton discovered the subject earlier, but it's his notation that we still use. He produced lasting treatises on metaphysics and theology, he drew up designs for watches and windmills, he advocated universal health care and the development of submarines. As a diplomat in Paris, he pressed Louis XIV to invade Egypt, a bank-shot ploy to divert Germany's mighty

neighbor into an overseas adventure that might lessen the prospect of marching its armies east. Denis Diderot, no slouch, moaned, "When one compares . . . one's own small talents with those of a Leibniz, one is tempted to throw away one's books and go die peacefully in the depths of some dark corner."

Of all Leibniz's schemes, the dearest was a new lexicon he called the universal characteristic—and it, too, sprang from his desire for peace. Throughout history, fanciful thinkers have created languages from scratch in the hope that their concoctions would smooth communication between the peoples of the world, fostering the preconditions for global oneness. Leibniz created his language for that reason, too, but he also had higher hopes: He argued that a new set of symbols and expressions would lead science and philosophy to new truths, to a new age of reason, to a deeper appreciation of the universe's elegance and harmony, to the divine.

What he imagined was an alphabet of human thought. It was an idea that he first pondered as a young student, the basis for his doctoral dissertation at Altdorf. Over the years, he fleshed out a detailed plan for realizing his fantasy. A group of scholars would create an encyclopedia containing the fundamental, incontestably true concepts of the world, of physics, philosophy, geometry, everything really. He called these core concepts "primitives," and they would include things like the earth, the color red, and God. Each of the primitives would be assigned a numerical value, which allowed them to be combined to create new concepts or to express complex extant ones. And those numerical values would form the basis for a new calculus of thought, what he called the calculus ratiocinator.

Leibniz illustrated his scheme with an example. What is a

human? A rational animal, of course. That's an insight that we can write like this:

$$\text{rational} \times \text{animal} = \text{man}$$

But Leibniz translated this expression into an even more mathematical sentence. "Animal," he suggested, might be represented with the number two; "rational" with the number three. Therefore:

$$2 \times 3 = 6$$

Thought had been turned into math—and this allowed for a new, foolproof method for adjudicating questions of truth. Leibniz asked, for instance, are all men monkeys? Well, he knew the number assigned to monkeys, ten. If ten can't be divided by six, and six can't be divided by ten, then we know: There's no element of monkey in man—and no element of man in monkey.

That was the point of his language: Knowledge, all knowledge, could ultimately be derived from computation. It would be an effortless process, *cogitatio caeca* or blind thought. Humans were no longer even needed to conceive new ideas. A machine could do that, by combining and dividing concepts. In fact, Leibniz built a prototype of such a machine, a gorgeous, intricate compilation of polished brass and steel, gears and dials. He called it the Stepped Reckoner. Leibniz spent a personal fortune building it. With a turn of the crank in one direction the Stepped Reckoner could multiply, in the other direction divide. Leibniz had designed a user interface so meticulous that Steve Jobs would have bowed down before it. Sadly, whenever he tested the machine

for an audience, as he did before the Royal Society in London in 1673, it failed. The resilient Leibniz forgave himself these humiliating demonstrations. The importance of the universal characteristic demanded that he press forward. "Once this has been done, if ever further controversies should arise, there should be no more reason for disputes between two philosophers than between two calculators." Intellectual and moral argument could be settled with the disagreeing parties declaring, "Let's calculate!" There would be no need for wars, let alone theological controversy, because truth would be placed on the terra firma of math.

Leibniz was a prophet of the digital age, though his pregnant ideas sat in the waiting room for centuries. He proposed a numeric system that used only zeros and ones, the very system of binary on which computing rests. He explained how automation of white-collar jobs would enhance productivity. But his critical insight was mechanical thinking, the automation of reason, the very thing that makes the Internet so miraculous, and the power of the tech companies so potentially menacing.

THOSE PROCEDURES THAT enable mechanical thinking came to have a name. They were dubbed algorithms. The essence of the algorithm is entirely uncomplicated. The textbooks compare them to recipes—a series of precise steps that can be followed mindlessly. This is different from equations, which have one correct result. Algorithms merely capture the process for solving a problem and say nothing about where those steps ultimately lead.

These recipes are the crucial building blocks of software. Programmers can't simply order a computer to, say, search the Internet. They must give the computer a set of specific instructions for

accomplishing that task. These instructions must take the messy human activity of looking for information and transpose that into an orderly process that can be expressed in code. First do this . . . then do that. . . . The process of translation, from concept to procedure to code, is inherently reductive. Complex processes must be subdivided into a series of binary choices. There's no equation to suggest a dress to wear, but an algorithm could easily be written for that—it will work its way through a series of either/or questions (morning or night, winter or summer, sun or rain), with each choice pushing to the next.

Mechanical thinking was exactly what Alan Turing first imagined as he collapsed on his run through the meadows of Cambridge in 1935 and daydreamed about a fantastical new calculating machine. For the first decades of computing, the term "algorithm" wasn't much mentioned. But as computer science departments began sprouting across campuses in the sixties, the term acquired a new cachet. Its vogue was the product of status anxiety. Programmers, especially in the academy, were anxious to show that they weren't mere technicians. They began to describe their work as algorithmic, in part because it tied them to one of the greatest of all mathematicians—the Persian polymath Muḥammad ibn Mūsā al-Khwārizmī, or as he was known in Latin, Algoritmi. During the twelfth century, translations of al-Khwārizmī introduced Arabic numerals to the West; his treatises pioneered algebra and trigonometry. By describing the algorithm as the fundamental element of programming, the computer scientists were attaching themselves to a grand history. It was a savvy piece of name dropping: See, we're not arriviste, we're working with abstractions and theories, just like the mathematicians!

There was sleight of hand in this self-portrayal. The algo-

rithm may be the essence of computer science—but it's not precisely a scientific concept. An algorithm is a system, like plumbing or a military chain of command. It takes know-how, calculation, and creativity to make a system work properly. But some systems, like some armies, are much more reliable than others. A system is a human artifact, not a mathematical truism. The origins of the algorithm are unmistakably human, but human fallibility isn't a quality that we associate with it. When algorithms reject a loan application or set the price for an airline flight, they seem impersonal and unbending. The algorithm is supposed to be devoid of bias, intuition, emotion, or forgiveness. They call it a search engine, after all—a nod to pistons, gears, and twentieth-century industry, with the machinery wiped clean of human fingerprints.

Silicon Valley's algorithmic enthusiasts were immodest about describing the revolutionary potential of their objects of affection. Algorithms were always interesting and valuable, but advances in computing made them infinitely more powerful. The big change was the cost of computing. It collapsed, and just as the machines themselves sped up and were tied into a global network. Computers could stockpile massive piles of unsorted data—and algorithms could attack this data to find patterns and connections that would escape human analysts. In the hands of Google and Facebook, these algorithms grew ever more powerful. As they went about their searches, they accumulated more and more data. Their machines assimilated all the lessons of past searches, using these learnings to more precisely deliver the desired results.

For the entirety of human existence, the creation of knowledge was a slog of trial and error. Humans would dream up theories of how the world worked, then would examine the evidence to see whether their hypotheses survived or crashed upon their

exposure to reality. Algorithms upend the scientific method—the patterns emerge from the data, from correlations, unguided by hypotheses. They remove humans from the whole process of inquiry. Writing in *Wired*, Chris Anderson argued: “We can stop looking for models. We can analyze the data without hypotheses about what it might show. We can throw the numbers into the biggest computing clusters the world has ever seen and let statistical algorithms find patterns where science cannot.”

On one level, this is undeniable. Algorithms can translate languages without understanding words, simply by uncovering the patterns that undergird the construction of sentences. They can find coincidences that humans might never even think to seek. Walmart’s algorithms found that people desperately buy strawberry Pop-Tarts as they prepare for massive storms. Still, even as an algorithm mindlessly implements its procedures—and even as it learns to see new patterns in the data—it reflects the minds of its creators, the motives of its trainers. Both Amazon and Netflix use algorithms to make recommendations about books and films. (One-third of purchases on Amazon come from these recommendations.) These algorithms seek to understand our tastes, and the tastes of like-minded consumers of culture. Yet the algorithms make fundamentally different recommendations. Amazon steers you to the sorts of books that you’ve seen before. Netflix directs users to the unfamiliar. There’s a business reason for this difference. Blockbuster movies cost Netflix more to stream. Greater profit arrives when you decide to watch more obscure fare. Computer scientists have an aphorism that describes how algorithms relentlessly hunt for patterns: They talk about torturing the data until it confesses. Yet this metaphor contains

unexamined implications. Data, like victims of torture, tells its interrogator what it wants to hear.

Sometimes, the algorithm reflects the subconscious of its creators. To take an extreme example: The Harvard professor Latanya Sweeney conducted a study that found that users with African American names were frequently targeted with Google ads that bluntly suggested that they had arrest records in need of expunging. (“Latisha Smith, Arrested?”) Google is not particularly forthright about why such results appear. Their algorithm is a ferociously guarded secret. Yet, we know that Google has explicitly built its search engine to reflect values that it holds dear. It believes that the popularity of a Web site gives a good sense of its utility; it chooses to suppress pornography in its search results and not, say, anti-Semitic conspiracists; it believes that users will benefit from finding recent articles more than golden oldies. These are legitimate choices—and perhaps wise business decisions—but they are choices, not science.

Like economics, computer science has its preferred models and implicit assumptions about the world. When programmers are taught algorithmic thinking, they are told to venerate efficiency as a paramount consideration. This is perfectly understandable. An algorithm with an ungainly number of steps will gum up the machinery, and a molasseslike server is a useless one. But efficiency is also a value. When we speed things up, we’re necessarily cutting corners, we’re generalizing.

Algorithms can be gorgeous expressions of logical thinking, not to mention a source of ease and wonder. They can track down copies of obscure nineteenth-century tomes in a few milliseconds; they put us in touch with long-lost elementary school friends;

they enable retailers to deliver packages to our doors in a flash. Very soon, they will guide self-driving cars and pinpoint cancers growing in our innards. But to do all these things, algorithms are constantly taking our measure. They make decisions about us and on our behalf. The problem is that when we outsource thinking to machines, we are really outsourcing thinking to the organizations that run the machines.

MARK ZUCKERBERG DISINGENUOUSLY POSES as a friendly critic of algorithms. That's how he implicitly contrasts Facebook with his rivals across the way at Google. Over in Larry Page's shop, the algorithm is king, a cold, pulseless ruler. There's not a trace of life force in its recommendations and very little apparent understanding of the person keying a query into its engine. Facebook, in his flattering self-portrait, is a respite from this increasingly automated, atomistic world. "Every product you use is better off with your friends," he says.

What he is referring to is Facebook's News Feed. Here's a brief explanation for the sliver of humanity who have apparently resisted Facebook: The News Feed provides a reverse chronological index of all the status updates, articles, and photos that your friends have posted to Facebook. The News Feed is meant to be fun, but also geared to solve one of the essential problems of modernity—our inability to sift through the ever-growing, always-looming mounds of information. Who better, the theory goes, to recommend what we should read and watch than our friends? Zuckerberg has boasted that the News Feed turned Facebook into a "personalized newspaper."

Unfortunately, our friends can do only so much to winnow

things for us. Turns out, they like to share a lot. If we just read their musings and followed links to articles, we might be only a little less overwhelmed than before, or perhaps even deeper underwater. So Facebook makes its own choices about what should be read. The company's algorithms sort the thousands of things a Facebook user could possibly see down to a smaller batch of choice items. And then within those few dozen items, it decides what we might like to read first.

Algorithms are, by definition, invisibilia. But we can usually sense their presence—that somewhere in the distance, we're interacting with a machine. That's what makes Facebook's algorithm so powerful. Many users—60 percent, according to the best research—are completely unaware of its existence. But even if they know of its influence, it wouldn't really matter. Facebook's algorithm couldn't be more opaque. When the company concedes its existence to reporters, it manages to further cloud the algorithm in impenetrable descriptions. We know, for instance, that its algorithm was once called EdgeRank. But Facebook no longer uses that term. It's appropriate that the algorithm doesn't have a name. It has grown into an almost unknowable tangle of sprawl. The algorithm interprets more than one hundred thousand "signals" to make its decisions about what users see. Some of these signals apply to all Facebook users; some reflect users' particular habits and the habits of their friends. Perhaps Facebook no longer fully understands its own tangle of algorithms—the code, all sixty million lines of it, is a palimpsest, where engineers add layer upon layer of new commands. (This is hardly a condition unique to Facebook. The Cornell University computer scientist Jon Kleinberg cowrote an essay that argued, "We have, perhaps for the first time ever, built machines we do not understand. . . . At

some deep level we don't even really understand how they're producing the behavior we observe. This is the essence of their incomprehensibility." What's striking is that the "we" in that sentence refers to the creators of code.)

Pondering the abstraction of this algorithm, imagine one of those earliest computers with its nervously blinking lights and long rows of dials. To tweak the algorithm, the engineers turn the knob a click or two. The engineers are constantly making small adjustments, here and there, so that the machine performs to their satisfaction. With even the gentlest caress of the metaphorical dial, Facebook changes what its users see and read. It can make our friends' photos more or less ubiquitous; it can punish posts filled with self-congratulatory musings and banish what it deems to be hoaxes; it can promote video rather than text; it can favor articles from the likes of the *New York Times* or BuzzFeed, if it so desires. Or if we want to be melodramatic about it, we could say Facebook is constantly tinkering with how its users view the world—always tinkering with the quality of news and opinion that it allows to break through the din, adjusting the quality of political and cultural discourse in order to hold the attention of users for a few more beats.

But how do the engineers know which dial to twist and how hard? There's a whole discipline, data science, to guide the writing and revision of algorithms. Facebook has a team, poached from academia, to conduct experiments on users. It's a statistician's sexiest dream—some of the largest data sets in human history, the ability to run trials on mathematically meaningful cohorts. When Cameron Marlow, the former head of Facebook's data science team, described the opportunity, he began twitching with ecstatic joy. "For the first time," Marlow said, "we have

a microscope that not only lets us examine social behavior at a very fine level that we've never been able to see before but allows us to run experiments that millions of users are exposed to."

Facebook likes to boast of the fact of its experimentation more than the details of the actual experiments themselves. But there are examples that have escaped the confines of its laboratories. We know, for example, that Facebook sought to discover whether emotions are contagious. To conduct this trial, Facebook attempted to manipulate the mental state of its users. For one group, Facebook excised the positive words from the posts in the News Feed; for another group, it removed the negative words. Each group, it concluded, wrote posts that echoed the mood of the posts it had reworded. This study was roundly condemned as invasive, but it is not so unusual. As one member of Facebook's data science team confessed: "Anyone on that team could run a test. They're always trying to alter people's behavior."

There's no doubting the emotional and psychological power possessed by Facebook—at least Facebook doesn't doubt it. It has bragged about how it increased voter turnout (and organ donation) by subtly amping up the social pressures that compel virtuous behavior. Facebook has even touted the results from these experiments in peer-reviewed journals: "It is possible that more of the .60% growth in turnout between 2006 and 2010 might have been caused by a single message on Facebook." No other company has so precisely boasted about its ability to shape democracy like this—and for good reason. It's too much power to entrust to a corporation.

The many Facebook experiments add up. The company believes that it has unlocked social psychology and acquired a deeper understanding of its users than they possess of themselves.

Facebook can predict users' race, sexual orientation, relationship status, and drug use on the basis of their "likes" alone. It's Zuckerberg's fantasy that this data might be analyzed to uncover the mother of all revelations, "a fundamental mathematical law underlying human social relationships that governs the balance of who and what we all care about." That is, of course, a goal in the distance. In the meantime, Facebook will probe—constantly testing to see what we crave and what we ignore, a never-ending campaign to improve Facebook's capacity to give us the things that we want and things that we don't even know we want. Whether the information is true or concocted, authoritative reporting or conspiratorial opinion, doesn't really seem to matter much to Facebook. The crowd gets what it wants and deserves.

THE AUTOMATION OF THINKING: We're in the earliest days of this revolution, of course. But we can see where it's heading. Algorithms have retired many of the bureaucratic, clerical duties once performed by humans—and they will soon begin to replace more creative tasks. At Netflix, algorithms suggest the genres of movies to commission. Some news wires use algorithms to write stories about crime, baseball games, and earthquakes, the most rote journalistic tasks. Algorithms have produced fine art and composed symphonic music, or at least approximations of them.

It's a terrifying trajectory, especially for those of us in these lines of work. If algorithms can replicate the process of creativity, then there's little reason to nurture human creativity. Why bother with the tortuous, inefficient process of writing or painting if a computer can produce something seemingly as good and in a painless flash? Why nurture the overinflated market for high cul-

ture, when it could be so abundant and cheap? No human endeavor has resisted automation, so why should creative endeavors be any different?

The engineering mind-set has little patience for the fetishization of words and images, for the mystique of art, for moral complexity and emotional expression. It views humans as data, components of systems, abstractions. That's why Facebook has so few qualms about performing rampant experiments on its users. The whole effort is to make human beings predictable—to anticipate their behavior, which makes them easier to manipulate. With this sort of cold-blooded thinking, so divorced from the contingency and mystery of human life, it's easy to see how long-standing values begin to seem like an annoyance—why a concept like privacy would carry so little weight in the engineer's calculus, why the inefficiencies of publishing and journalism seem so imminently disruptable.

Facebook would never put it this way, but algorithms are meant to erode free will, to relieve humans of the burden of choosing, to nudge them in the right direction. Algorithms fuel a sense of omnipotence, the condescending belief that our behavior can be altered, without our even being aware of the hand guiding us, in a superior direction. That's always been a danger of the engineering mind-set, as it moves beyond its roots in building inanimate stuff and begins to design a more perfect social world. We are the screws and rivets in the grand design.