

pulling me down. Then I told myself: You're making excuses for yourself. You are doing this, nobody else. These are *your* failings. And I felt weak. I had gained a lot of insights in Provincetown—but I felt they were fragile, and easily broken by something bigger, something I didn't quite yet understand.

I wanted to know what was preventing me from doing what I wanted, in the better part of myself, to do. I discovered that the answer is more complex than we have been led to believe, and has many facets—and I learned about the first of them when I went to Silicon Valley.

CHAPTER SIX

Cause Six: The Rise of Technology That Can Track and Manipulate You (Part One)

James Williams told me I had made a fundamental mistake in Provincetown. He was a senior Google strategist for many years, and he left, horrified, to go to Oxford University, to study human attention, and figure out what his colleagues in Silicon Valley have done to it. He told me a digital detox is “not the solution, for the same reason that wearing a gas mask for two days a week outside isn't the answer to pollution. It might, for a short period of time, keep, at an individual level, certain effects at bay. But it's not sustainable, and it doesn't address the systemic issues.” He said our attention is being deeply altered by huge invasive forces in the wider society. Saying the solution is primarily to personally abstain is just “pushing it back onto the individual,” he said, when “it's really the environmental changes that will really make the difference.”

For a long time I didn't fully understand what this meant. What would changing our environment entail, when it came to attention, if not each of us trying to change our own personal behavior? The

answer slowly became clear to me when I met with many people who had designed crucial aspects of the world in which we now live. In the hills of San Francisco and the hot, arid streets of Palo Alto, I realized that there are six ways in which our technology, as it currently works, is harming our ability to pay attention—and that these causes are united by one deeper underlying force that needs to be overcome.

One of the first people to guide me on this journey was Tristan Harris, another former Google engineer, who, after I had been interviewing him for several years, became globally famous for appearing in the viral Netflix documentary *The Social Dilemma*. That film explored a whole range of ways in which social media, as it is currently designed, can be destructive. I wanted to tease out something the film largely didn't explore—its effect on our focus. To grasp it, I think it helps to know Tristan's own story, and what he witnessed at the heart of the machine that is repatterning the world's attention.



In the early 1990s, in the town of Santa Rosa, California, a little boy with a bowl haircut and a bright golden bow tie was learning magic. Tristan was seven years old when he first tried out one of the most basic tricks. He would ask you to hand him a coin, and then—poof! It was gone. After he mastered more tricks, he put on a magic show for his elementary-school class, and then—to his glee—he was selected to go to a magic camp out in the hills, where he was taught for a week by professional magicians. It seemed to him like a real-life Jedi training camp.

He discovered, at this young age, the most important fact about magic. He explained years later: "It's really about the limits of attention." The job of a magician is—at heart—to manipulate your focus. That coin didn't really vanish—but your attention was somewhere else when the magician moved it, so when your focus comes back to the original spot, you're amazed. To learn magic is to learn to manipulate someone's attention without them even realizing it—and once

the magician controls their focus, Tristan realized, he can do what he wants. One of the things that he was taught at camp is that a person's susceptibility to magic has nothing to do with their intelligence. "It's about something more subtle," he said later. It's "about the weaknesses, or the limits, or the blind spots, or the biases that we're all trapped inside of."

Magic, in other words, is the study of the limits of the human mind. You think you control your attention; you think that if somebody messes with it, you will know, and you'll be able to spot and resist it right away, but, in reality, we are fallible sacks of meat, and we are fallible in predictable ways that can be figured out by magicians and messed with.

As he got to know better and better magicians—eventually befriending one of the best in the world, Derren Brown—Tristan learned something he found both remarkable and disconcerting. It is possible to manipulate your attention to such a degree that a magician can, in many cases, turn you into his puppet. He can make you choose whatever he wants you to choose, while all along you think you're simply using your own free will. When Tristan first said this to me, I thought he was overstating his case, so he introduced me to another magician friend, James Brown. Tristan told me James would show me what it meant. I'll give you one example. When we sat together, James showed me a standard pack of cards. He said: See? Some of them are red, and some of them are black, and they are all mixed up together. Then he turned the cards so the colors were facing toward him, and I couldn't see them anymore. He told me he was going to get me to sort them neatly into two piles—one black, one red—without me ever getting to look at the color of the cards for myself. It was, obviously, impossible. How could I sort cards I couldn't see?

He told me to look into his eyes, and—entirely using my own free will—to tell him whether to put the next card into a pile on the left, or a pile on the right. So I gave him my orders—left, left, right, and so on—according to what I was confident were my own random

whims. At the end, he lifted up the piles of cards and showed them to me. The red cards were neatly in one pile; the black cards were in the other.

I was baffled. How did he do it? He eventually told me he had been subtly guiding my choices. He said he would do it again, a little more crudely this time, to see if I could spot it. Finally—and he had to be pretty blatant—I saw it. When he told me to pick the next card, he indicated very slightly with his eyes to the left or to the right—and I always chose in the way he unconsciously guided me to. Everyone always does, he told me. Later, Tristan explained to me that this is a core insight of magic—you can manipulate people and they don't even know it's happening. They will swear to you that they made their own free choices—as I would have about those cards.

One morning, in his office in San Francisco, Tristan leaned forward and said to me: “How does a magician do their work? It works because they don't have to know your strengths—they just have to know your weaknesses. How well do you know your weaknesses?” I wanted to believe I understood my weaknesses very well, but Tristan shook his head gently. “If people did know their weaknesses,” he said, “then magic wouldn't work.”

Magicians play on these weaknesses to delight and entertain us. As Tristan grew up, he became part of another group of people who were figuring out our weaknesses to manipulate us—but they had very different goals.

It was in his first year at Stanford University, in 2002, that Tristan first heard whispers of a mysterious-sounding research facility known as the Persuasive Technology Lab. It was, the rumors went, a place where a leading scientist was figuring out how to design technology that could change your behavior in unprecedented ways. In his teens Tristan had become obsessed with coding, and he had already been an intern at Apple after his freshman year at Stanford, designing a piece of code that is still used in many of your devices today.

This much-discussed course, he learned, was about taking a huge amount of what scientists had discovered over the twentieth century about how to change other people's behavior, and enabling students to understand how these classic influence techniques could be used in tech products.

The course was taught by a warm, upbeat behavioral scientist in his forties named Professor B. J. Fogg. At the start of each day, he would take out a stuffed frog and a cuddly monkey and introduce them to the class, and then he would play on his ukulele. Whenever he wanted the group to break or wrap up, he would tap on a toy xylophone. B.J. explained to students that computers had the potential to be far more persuasive than people. They could, he believed, “be more persistent than human beings, offer greater anonymity,” and “go where humans cannot go or may not be welcome.” Soon, he was sure, they would be changing the lives of everyone—persuading us persistently, throughout the day. He had previously worked on a course dedicated to “the psychology of mind control.” In the fall of 2006, B.J. testified before the Federal Trade Commission (FTC), warning that there was a “potential dark side” looming, where it will be hard for you to tell “when you are being persuaded and when you are not.” He assigned to Tristan and his other students a large amount of books that explained dozens of psychological insights and tricks that had been discovered about how to change human beings and to get them to do what you want. It was a treasure trove. One of the texts that most struck Tristan was based on the philosophy of B. F. Skinner, the man who, as I had learned earlier, had found a way to get pigeons and rats and pigs to do whatever he wanted by offering the right “reinforcements” for their behavior. After falling out of fashion for years, his ideas were back with full force.

“It really woke up the magic part of me,” Tristan told me. “I was like—oh wow, there really are these invisible rules that govern what people do. And if there are rules that govern what people do, that's power. That's like Isaac Newton discovering the laws of physics. It felt like somebody's showing me the code—the code of how you can

influence people. I remember the experience of sitting there in the graduate area of campus reading those books over the weekends, and underlining furiously these passages, and just being like—oh my God, I can't even believe that works." He was so intoxicated by the excitement of it that, he said, "I will admit, I don't think the ethical bells were firing in my brain yet."

As part of the class, he was paired with a young man named Mike Krieger, and they were tasked with designing an app. Tristan had been thinking for a while about something named "seasonal affective disorder"—a condition where, if you are stuck in gloomy weather for a long time, you are more likely to become depressed. How, they asked, could technology help with that? They came up with an app called Send the Sunshine. Two friends would choose to be connected through it, and it would track where they both were and the online weather reports for their locations. If the app realized that your friend was starved of sunshine, and you had some, it would prompt you to take a photo of the sun and send it to him. It showed that somebody cared; and it sent some sunshine your way. It was sweet, and simple, and it helped to spur Mike to think about the power of sharing photographs online. He was already thinking about another of the key lessons of the class, taken from B. F. Skinner: build in immediate reinforcements. If you want to shape the user's behavior, make sure he gets hearts and likes right away. Using these principles, he and another friend, Kevin Systrom, launched a new app of their own. They named it Instagram.

The class was filled with people who were going to use the techniques B.J. taught to change how we live our lives, and B.J. was dubbed "the millionaire maker." But something was starting to nag at Tristan. After a while, he noticed he had become obsessed with checking his email. He would do it repetitively, mindlessly, again and again, and he felt his attention span was beginning to atrophy. He realized, he told me, that the email app he was using "operates on a bunch of different levers, and it's very powerful, and it sucks, and it's super-stressful, and it ruins hours and hours of people's lives." He had learned at Stanford and elsewhere how to use technology to

change people's behavior, but he came to ask a disconcerting question: Am I somehow being hacked by other tech designers myself? He wasn't yet sure how they might be doing it—but he began to have a strange feeling about it. B.J. taught his students that they should only use these powers for good, and he laced ethical debates throughout his course. Tristan told me: "B.J. is a good and ethical person who shouldn't bear the blame for what we're seeing take place." Yet Tristan was going to start to wonder: Were these radical new technologies actually being used ethically in the real world, by people who were leaving B.J.'s moral ambit, or had never heard his ethics lessons in the first place?

In the final class Tristan attended, all the students discussed ways in which these persuasive technologies could be used in the future. Professor Fogg led his students in a discussion, posing the question: "What if in the future you had a profile of every single person on earth?" As a designer, you would track all the information they offer up on social media and build a detailed profile of them. It's not just the simple stuff—their gender, or age, or interests. It would be something deeper. This would be a psychological profile—figuring out how their personality works, and the best ways to persuade them. It would know if the user was an optimist or pessimist, if they were open to new experiences or they were prone to nostalgia—it would figure out dozens of characteristics they have.

Think, the class wondered out loud, about how you could target people if you knew this much about them. Think about how you could change them. When a politician or a company wants to persuade you, they could pay a social-media company to perfectly target their message just for you. It was the birth of an idea—one that would be picked up in Silicon Valley later by different people far removed from B.J. and his warnings about ethics. Years later, when it was revealed that the campaign for Donald Trump had paid a company named Cambridge Analytica to do exactly that, Tristan would think of that final class discussion at Stanford. "This was the class that freaked me out," he told me. "I remember saying—this is horribly concerning."

But Tristan had a deep belief in the power of tech to do good. So he took what he had learned at Stanford and—under the mentorship of B.J., his former teacher at Stanford—designed an app with a straightforward positive purpose. He was trying to stop one of the ways the web screws with our attention. Let's say you are checking out the CNN site, and you start to read a news story about Northern Ireland, a topic you don't know much about. Normally, you will then open a new window and begin googling for info—and before you know it, you vanish down a rabbit hole and emerge half an hour later, lost in articles and videos about a totally different topic (usually cats playing the piano). Tristan's app was designed so that in this situation, you could do something different: you could highlight any phrase (say, "Northern Ireland"), and it would pull up a simple pop-up window giving you a straightforward summary of the topic. No clicking away from the site; no rabbit holes. Your attention is preserved. The app did well—it started to be used by thousands of websites, including the *New York Times*, and quite soon, Google made a substantial offer to buy the whole thing and for Tristan to come and work for them. They told him it was so he could integrate it into their web browser, Chrome, and make people less distracted. He jumped at the chance.

It is hard to convey, Tristan believes, quite what it was like to go to work for Google at that moment in history, in 2011. Every day, the company he worked for—from its base, the Googleplex in Palo Alto—was shaping and reshaping how one billion people navigated their way through the world: what they got to see, and what they didn't. He told one audience later: "I want you to imagine walking into a room. A control room, with a bunch of people, a hundred people, hunched over a desk with little dials—and that that control room will shape the thoughts and feelings of a billion people. This might sound like science fiction, but this actually exists right now, today. I know, because I used to work in one of those control rooms."

Tristan was assigned for a while to work on the development of

Gmail, Google's email system—precisely the app that was driving him wild, and that he suspected might be using some manipulative tricks he hadn't yet figured out. Even as he worked on it, he would obsessively check his email, making him less focused, and whenever he looked at a new message, he found it took him a long time to get his mind back to where it had been before. He started trying to think through how you might design a system of email that was less prone to nuking your attention—but whenever he tried to discuss this idea with his colleagues, the conversation didn't seem to go far. At Google, he quickly learned, success was measured, in the main, by what was called "engagement"—which was defined as minutes and hours of eyeballs on the product. More engagement was good; less engagement was bad. This was for a simple reason. The longer you make people look at their phones, the more advertising they see—and therefore the more money Google gets. Tristan's co-workers were decent people, struggling with their own tech distractions—but the incentives seemed to lead only one way: you should always design products that "engage" the maximum number of people, because engagement equals more dollars, and disengagement equals fewer dollars.

With each month that passed, Tristan became more startled by the casualness with which the attention of a billion people was being corroded at Google and the other Big Tech companies. One day he would hear an engineer excitedly saying: "Why don't we make it buzz your phone every time we get an email?" Everyone would be thrilled—and a few weeks later, all over the world, phones began to buzz in pockets, and more people found themselves looking at Gmail more times a day. The engineers were always looking for new ways to suck eyeballs onto their program and keep them there. Day after day, he would watch as engineers proposed more interruptions to people's lives—more vibrations, more alerts, more tricks—and they would be congratulated.

As the number of people using Google and Gmail continued to spike up, Tristan started to ask his colleagues: "How do you ethically persuade two billion people's minds? How do you ethically

structure two billion people's attention?" But instead, he found that most other people in the company were being pushed to ask simply, "How can we make this more engaging?" And that meant more attention-sucking, more interrupting—on and on it went, with better techniques being discovered every week. One day, when we were walking in San Francisco, Tristan said to me: "Things look pretty bad from the outside, but when you're on the inside, things can look even worse." Tristan was starting to realize: It's not your fault you can't focus. It's by design. Your distraction is their fuel.

After working intensively on the Gmail team, Tristan saw that when it came to questioning what they were doing to people's attention, "the conversation was not happening." He looked out across his friends now working in every part of Silicon Valley, and this grab-and-raid approach to our focus was being taken in almost all the companies they worked in. "What started to really concern me over the years," he told me, "was just watching my friends who had originally gotten into this business because they thought they could make the world better, [and now] were caught in this arms race to manipulate human nature."

To pluck one example out of dozens Tristan could offer, his friends Mike and Kevin had launched Instagram, and after a little while, "they added these filters, because it was a cool thing. So you could take a photo, and just have it look artistic instantly." It didn't cross their minds, he's sure, that it would start a race with Snapchat and others to see who could "provide better beautification filters"—and that this would, in turn, change how people thought of their own bodies so much that today there's a whole category of people who undergo surgery so they can look more like their filters. He could see that his friends were setting in motion changes that were transforming the world in ways they couldn't predict or control. "The reason we have to be so careful about the way that we design technology," he said, is that "they squeeze, they squish, the entire world down into that medium—and out the other end comes a different world."

But here was Tristan, at the center of the machine unleashing

these transformations, and he could see that behind closed doors, the dials in the control room were being set to ten.



After a few years at the heart of the Googleplex, Tristan couldn't take it anymore, and he decided to leave. As a final gesture, he put together a slide-deck for the people he worked with, to appeal to them to think about these questions. The first slide said simply: "I'm concerned about how we're making the world more distracted." He explained: "Distraction matters to me, because time is all we have in life. . . . Yet hours and hours can get mysteriously lost here." He showed a picture of a Gmail inbox. "And [on] feeds that suck huge chunks of time away here." He showed a Facebook feed. He said he was worried that the company—and others like it—was inadvertently "destroy[ing] our kids' ability to focus," pointing out that the average child between the ages of thirteen and seventeen in the U.S. was sending one text message every six minutes they were awake. People were, he warned, living "on a treadmill of continuous checking."

He asked: We know that interruptions cause a deterioration in people's ability to focus and think clearly—so why are we ramping up the interruptions? Why are we finding better and better ways to do it all the time? "Think about that," he told his colleagues. "We should feel an enormous responsibility to get this right." All humans have natural vulnerabilities, and instead of exploiting those vulnerabilities—like a malign magician—Google should be respecting them. He suggested some modest changes as a place to start. Instead of notifying someone every time they have a new email, he suggested, we could notify them once a day, in a batch—so it'd be like getting a newspaper in the morning, instead of constantly following the rolling news. Every time we prompt somebody to click over to a new photo their friend has posted, we could warn them—on the same screen—that the average person who clicks on a photo is pulled away for twenty minutes before they get back to their task. We could tell them: You think it'll only take a second, but it won't.

He suggested giving users a chance to pause every time they click to do something potentially seriously distracting, to check: Are you sure you want to do this? Do you know how much time it will take from you? "Humans make different decisions when we pause and consider," he said.

He was trying to give his colleagues a sense of the weight of the decisions they made every day: "We shape more than eleven billion interruptions to people's lives *every day*. This is nuts!" The people sitting around you in the Googleplex, he explained, control more than 50 percent of all the notifications on all the phones in the whole world. We are "creating an arms race that causes companies to find more reasons to steal people's time," and it "destroys our common silence and ability to think." He asked: "Do we really know what we're doing to people?"

This was an almost insanely bold thing to do. At the heart of the machine that was changing the world, here was a smart and talented but fairly junior engineer, still only twenty-nine years old, saying something that directly challenged the whole direction of the company. It would be like a junior exec in 1975 standing up in front of the whole of ExxonMobil and telling them that they were responsible for global warming by showing them images of the melting of the Arctic. Everyone in Silicon Valley was scrambling to get into and suck up to Google. But here was Tristan, with the ability to stay at its heart forever and make a lot of money, writing what seemed to be his own professional death certificate, because he believed somebody, somewhere, had to say something.

He shared his slide-deck with his colleagues, and went home, depressed. Then something unexpected happened.

With each hour that passed, more and more Google employees shared Tristan's slide-deck. The next day, he was inundated with messages from within the company enthusing about it. It turned out he had

tapped into a latent mood. Just because you design these products, it doesn't mean you are more insulated than anyone else from becoming hooked on them. The workers at the Googleplex could feel this tsunami of distractions hitting them too. Many of them wanted to have a serious conversation about what they were doing to the world. People were drawn in particular to the question Tristan had put to them: "What if we designed [our products] to minimize stress and create calmer states of mind?"

There was some pushback too. A few of his colleagues said that every new technology brings with it a panic where people say it'll trash the world—after all, Socrates said writing things down would ruin people's memories. We were told that everything from the printed book to television would trash the minds of the young, but here we are, and the world survived. Some others responded from a libertarian perspective, saying that what he was suggesting would invite government regulation, which they believed was contrary to the whole spirit of cyberspace.

Tristan's presentation caused such a ruction within Google that he was asked to stay in a special new position, created just for him. They offered him the role of being Google's first "design ethicist." He was thrilled. Here was a chance to think through some of the most challenging questions of our time, in a place where—if he could get people to listen—he could make an enormous difference. For the first time in a long time, he felt optimistic. He thought his new appointment meant Google was serious about exploring these questions. He knew there was enthusiasm for it among his fellow workers, and he believed in the good faith of his bosses.

He was assigned a desk, and—in effect—left to think. So he started to research the effects of many things. For example, he looked at the way Snapchat hooks teenagers. The app had an option called "Snapchat streaks," where two friends—almost always teens—would check in with each other every day through the app. Every day they checked in, their streak got longer, so you would aim to build up a

streak of two hundred, three hundred, four hundred days, all on a brightly colored display full of emojis. If you missed a single day, it would reset to zero. It was a perfect way to take the desire of teens for social connection and manipulate it to get them hooked. You came every day to extend your streak, and you stuck around to scroll, often for hours.

But whenever he came up with a specific proposal for how Google's own products could be less interrupting and presented it to people above him, he was told, in effect: "This is hard, it's confusing, and it's often at odds with our bottom line." Tristan realized he was bumping up against a core contradiction. The more people stared at their phones, the more money these companies made. Period. The people in Silicon Valley did not want to design gadgets and websites that would dissolve people's attention spans. They're not the Joker, trying to sow chaos and make us dumb. They spend a lot of their own time meditating and doing yoga. They often ban their own kids from using the sites and gadgets they design, and send them instead to tech-free Montessori schools. But their business model can only succeed if they take steps to dominate the attention spans of the wider society. It's not their goal, any more than ExxonMobil deliberately wants to melt the Arctic. But it's an inescapable effect of their current business model.

When Tristan warned about these negative effects, most people inside the company sympathized and agreed. When he suggested alternatives, people changed the subject. To give you a sense of the money involved: the personal wealth of Larry Page, one of the founders of Google, is \$102 billion; his colleague Sergey Brin is worth \$99 billion; and their colleague Eric Schmidt is worth \$20.7 billion. That's separate from Google's wealth as a company, which as I write stands at \$1 trillion. These three men alone are worth roughly the same as the total combined wealth of every single person, building, and bank account in the oil-rich country of Kuwait, and Google is worth roughly the entire wealth of the whole of Mexico or Indonesia. Telling them to distract people less was like telling an oil company

not to drill for oil—they didn't want to hear it. "You don't even really get to make that ethical decision" to improve people's attention spans, Tristan realized, "because your business model and your incentives are making that decision for you." Years later, testifying before the U.S. Senate, he explained: "I failed because companies don't [currently] have the right incentive to change."

Tristan was in the ethicist job for two years, and toward the end, as he told an audience later, "I felt completely hopeless. There were literally days when I went to work and I would read Wikipedia all day and check my email and I would have no idea how, once you see something as massive as the attention economy and its perverse incentives, a system this big could ever change. I truly felt hopeless. I felt depressed." So, finally, he quit Google, and went out into a Silicon Valley where, as he put it to me, "everything is a race for attention." In that lonely time in Tristan's life, he was about to team up with another person who felt depressed and lost—and who felt guilty about what he personally had done to you, me, and everyone we know.

You probably haven't heard of Aza Raskin, but he has directly intervened in your life. He will, in fact, probably affect how you spend your time today. Aza grew up in the most elite sliver of Silicon Valley, at the height of its confidence that it was making the world better. His dad was Jef Raskin, the man who invented the Apple Macintosh for Steve Jobs, and he built it around one core principle: that the user's attention is sacred. The job of technology, Jef believed, was to lift people up and make it possible to achieve their higher goals. He taught his son: "What is technology for? Why do we even make technology? We make technology because it takes the parts of us that are most human and it extends them. That's what a paintbrush is. That's what a cello is. That's what language is. These are technologies that extend some part of us. Technology is not about making us super-human. It's about making us extra-human."

Aza became a precocious young coder, and he gave his first talk about user interfaces when he was ten years old. By the time he was in his early twenties, he was at the forefront of designing some of the first internet browsers, and he was the creative lead on Firefox. As part of this work, he designed something that distinctly changed how the web works. It's called "infinite scroll." Older readers will remember that it used to be that the internet was divided into pages, and when you got to the bottom of one page, you had to decide to click a button to get to the next page. It was an active choice. It gave you a moment to pause and ask: Do I want to carry on looking at this? Aza designed the code that means you don't have to ask that question anymore. Imagine you open Facebook. It downloads a chunk of status updates for you to read through. You scroll down through it, flicking your finger—and when you get to the bottom, it will automatically load another chunk for you to flick through. When you get to the bottom of that, it will automatically load another chunk, and another, and another, forever. You can never exhaust it. It will scroll infinitely.

Aza was proud of the design. "At the outset, it looks like a really good invention," he told me. He believed he was making life easier for everyone. He had been taught that increased speed and efficiency of access were always advances. His invention quickly spread all over the internet. Today, all social media and lots of other sites use a version of infinite scroll. But then Aza watched as the people around him changed. They seemed to be unable to pull themselves away from their devices, flicking through and through and through, thanks in part to the code he had designed. He found himself infinitely scrolling through what he often realized afterward was crap, and he wondered if he was making good use of his life.

One day, when he was thirty-two, Aza sat down and did a calculation. At a conservative estimate, infinite scroll makes you spend 50 percent more of your time on sites like Twitter. (For many people, Aza believes, it's vastly more.) Sticking with this low-ball percentage, Aza wanted to know what it meant, in practice, if billions of people

were spending 50 percent more on a string of social-media sites. When he was done, he stared at the sums. Every day, as a direct result of his invention, the combined total of 200,000 more human lifetimes—every moment from birth to death—is now spent scrolling through a screen. These hours would otherwise have been spent on some other activity.

When he described this to me, he still sounded a little stunned. That time is "just completely gone. It's like their entire life—poof. That time, which could have been used for solving climate change, for spending time with their family, for strengthening social bonds. For whatever it is that makes their life well-lived. It just . . ." He trailed off. I pictured my young godson Adam and all his teenage friends, scrolling, scrolling, infinitely scrolling.

Aza told me he felt "sort of dirty." He realized: "These things we do, they really can change the world. Then the question immediately follows: In what way did we change the world?" He realized he thought making tech easier to use meant the world would get better. But he began to think that "one of my biggest learnings as a designer or technologist is—making something easy to use doesn't mean it's good for humanity." He thought about his father—who had since died—and his commitment to make tech that set people free to be better, and he wondered if he was living up to his dad's vision. He began to ask if he and his generation in Silicon Valley were actually "mak[ing] technology that tears us, rips us, and breaks us."

He carried on designing more things in the vein of infinite scroll, and getting more and more uncomfortable. "It was about the time that we were getting to be really successful at this that my stomach started to drop," he told me. He felt that he was seeing people become more unempathetic, angry, and hostile as their social-media use went up. At the time, he was running an app he had designed named Post-Social, which was a social-media site designed to help people interact more in the real world, away from their devices. He was trying to raise money for the next phase of its development, and all any investor wanted to know was: How much of people's

attention do you capture and run through your app? How often? How many times a day? That's not what Aza wanted to be—a person who thought solely about how to drain away people's time. But "you could see this gravity, pulling this product back to everything that we were trying to fight against."

The logic of the underlying system was being laid bare for Aza. Silicon Valley sells itself by articulating "a big, lofty goal—connecting everyone in the world, or whatever it is. But when you're actually doing the day-to-day work, it's about increasing user numbers." What you are selling is your ability to grab and hold attention. When he tried to discuss this, he thwacked into raw denial. "Say you were baking bread," he said to me, "and you had this incredible bread, and you used this secret substance—and all of a sudden, you're making free bread for the world, and everyone's eating it. Then one of your scientists comes and says: 'By the way, we think it causes cancer, this secret substance.' What do you do? You would almost certainly say: 'That can't be right. We need more research. Maybe it's something [else] that the people out there are doing. Maybe there's some other factor.'"

All throughout the industry, Aza kept meeting people who were going through similar crises. "There were a number of dark nights of the soul that I personally witnessed," he says. He watched as Silicon Valley's own inhabitants seemed to be hijacked by their own creations, and then tried to escape. When I met with several of these tech dissidents, it struck me how young they were—like they were almost children who had invented toys and watched their toys conquer the world. Everyone was scrambling to meditate in an attempt to resist the programs they had invented. He realized "one of the ironies is there are these incredibly popular workshops at Facebook and Google about mindfulness—about creating the mental space to make decisions nonreactively—and they are also the biggest perpetrators of non-mindfulness in the world."

When Tristan and Aza started to speak out, they were ridiculed as wildly over-the-top Cassandras. But then, one by one, all over Silicon Valley, people who had built the world we now live in were beginning to declare in public that they had similar feelings. For example, Sean Parker, one of the earliest investors in Facebook, told a public audience that the creators of the site had asked themselves from the start: "How do we consume as much of your time and conscious attention as possible?" The techniques they used were "exactly the kind of thing that a hacker like myself would come up with, because you're exploiting a vulnerability in human psychology. . . . The inventors, creators—it's me, it's Mark [Zuckerberg], it's Kevin Systrom on Instagram, it's all of these people—understood this consciously. And we did it anyway." He added: "God only knows what it's doing to our children's brains." Chamath Palihapitiya, who had been Facebook's vice president of growth, explained in a speech that the effects are so negative that his own kids "aren't allowed to use that shit." Tony Fadell, who co-invented the iPhone, said: "I wake up in cold sweats every so often thinking, what did we bring to the world?" He worried that he had helped create "a nuclear bomb" that can "blow up people's brains and reprogram them."

Many Silicon Valley insiders predicted that it would only get worse. One of its most famous investors, Paul Graham, wrote: "Unless the forms of technological progress that produced these things are subject to different laws than technological progress in general, the world will get more addictive in the next forty years than it did in the last forty."

One day, James Williams—the former Google strategist I met—addressed an audience of hundreds of leading tech designers and asked them a simple question: "How many of you want to live in the world you are designing?" There was a silence in the room. People looked around them. Nobody put up their hand.

CHAPTER NINE

The First Glimpses of the Deeper Solution

After learning so much about how our technology works, I was left with two clear and urgent questions. Firstly: what are the specific changes to this invasive tech that could be made, in practice, to prevent it from harming our attention and focus? And secondly: how do we compel these huge corporations to introduce these changes in the real world?

Tristan and Aza—drawing on their own experiences, and the essential work of Professor Shoshana Zuboff—believe that if we are going to find a lasting solution, we need to go right to the root cause of the problem. That’s why, one morning, Aza said to me starkly: “We could just ban surveillance capitalism.” I paused to try to process what he was saying. This would mean, he explained, that the government would ban any business model that tracks you online in order to figure out your weaknesses and then sells that private data to the highest bidder so they can change your behavior. This model is, Aza

says, “just fundamentally anti-democratic and anti-human,” and it has to go.

This sounded dramatic and frankly impossible when I first heard it, but Tristan and Aza explained that there are plenty of historical precedents for something becoming so widespread, only for society to discover it actually causes a lot of harm and ban the market in it. Think about lead paint. It was in a majority of American homes—then it was discovered to damage the brains of kids and adults, making it harder for them to focus. As one of Tristan’s mentors, Jaron Lanier, pointed out to me, when we found that out, we didn’t say nobody could paint their homes ever again. We just banned the lead in the paint. Your home is still painted today—just with much better products. Or think about CFCs. As I mentioned before, when I was a child in the hairspray-obsessed 1980s, it was discovered that a substance in hairsprays was destroying the ozone layer that protects us from the sun’s rays. It terrified us all. We banned CFCs. We still have hairsprays, they just work differently, and today, the ozone layer is healing. There are all sorts of things that we have decided, as a civilized society, can’t be bought and sold, like (for example) human organs.

So, I asked them, let’s say we banned surveillance capitalism. What would happen to my Facebook and Twitter accounts the following day, the following week, the following year? “I think they would have a crisis moment, in the same way that Microsoft had a crisis moment,” Aza told me. In 2001, Microsoft was ruled by the U.S. government to have become a monopoly. That company reinvented itself, and now “they’re sort of like the benevolent adult in the room. I think the same transformation would happen to Facebook.”

In practice, the day after a ban, these companies would have to find different ways to fund themselves. One model that is obvious—and an alternative form of capitalism that everyone reading this will have some experience of—is subscription. Let’s imagine each of us had to pay fifty cents or a dollar every month to use Facebook.

Suddenly, Facebook would no longer be working for advertisers and offering up your secret wishes and preferences as their real product. No. It would be working for you. Its job, for the first time, would be to actually figure out what makes you happy, and to give it to you—instead of figuring out what makes advertisers happy, and how they can manipulate you to give it to them. So if, like most people, you want to be able to focus, the site would have to be redesigned to facilitate that. If you want to be socially connected, instead of isolated in front of your screen, it would have to figure out how to make that possible.

There's another obvious way that these companies could survive, which is for them to be bought by the government and taken into public ownership. This would take social media out of the capitalist part of the economy. This can sound drastic, but every single person reading this book benefits today, directly, from exactly the same model. We all agree we need to have sewers—they are an unavoidable necessity, unless we want to go back to the world of cholera outbreaks and feces in the streets. So in virtually every country, the government owns, maintains, and regulates the sewers, and even hardcore anti-government activists agree that this is a good use of state power.

Using the same model, our governments could acknowledge that social media is now an essential public utility, and explain that when it is run according to the wrong incentives, it causes the psychological equivalents of cholera outbreaks. It would be a bad idea for the government to run it—it's easy to imagine how authoritarian leaders could abuse that. Fortunately, there's a better option: you can have public ownership, independent of the government. In Britain, the BBC is owned and funded by the British public, and it is run in the interests of the British public—but its day-to-day running is independent of the government. It's not perfect, but this model works so well that it is the most respected media organization in the world.

Once the financial incentives are changed—either through subscription or public ownership, or another model—then the nature

of these sites can change, in ways we can actually begin to envision already. Aza told me that "it's actually technically not hard" to redesign the major social-media sites so that, instead of trashing your attention span and our societies, they would be designed to heal them, once the financial incentives to do that are in place. This was difficult for me to grasp at first, so I asked what social media would look like after the changes they would like to see. Tristan, Aza, and others began by explaining minor changes, then built up to big changes, then told me what has to occur to make any of these changes happen.

They started by talking about how these companies could, overnight, remove a lot of the aspects of these apps and sites that deliberately scramble our heads and keep us online longer than we really want. Aza said: "For instance, Facebook tomorrow could start batching your notifications, so you only get one push notification a day. . . . They could do that tomorrow." (This was something Tristan had proposed in his explosive slide-deck back when he was still at Google.) So instead of getting "this constant drip of behavioral cocaine," telling you every few minutes that somebody liked your picture, commented on your post, has a birthday tomorrow, and on and on and on—you would get one daily update, like a newspaper, summarizing it all. You'd be pushed to look once a day, instead of being interrupted several times an hour.

"Here's another one," he said. "Infinite scroll." That's his invention, where when you get to the bottom of the screen, it automatically loads more and more, forever. "What's going on there is it's catching your impulses before your brain has a chance to really get involved and make a decision." Facebook and Instagram and the others could simply turn off infinite scroll—so that when you get to the bottom of the screen, you have to make a conscious decision to carry on scrolling.

Similarly, these sites could simply switch off the things that have been shown to most polarize people politically, stealing our ability to pay collective attention. Since there's evidence YouTube's

recommendation engine is radicalizing people, Tristan told one interviewer: "Just turn it off. They can turn it off in a heartbeat." It's not as if, he points out, the day before recommendations were introduced, people were lost and clamoring for somebody to tell them what to watch next.

Once the most obvious forms of mental pollution have been stopped, they said, we can begin to look deeper, at how these sites could be redesigned to make it easier for you to restrain yourself and think about your longer-term goals. "It doesn't take much work to start imagining what would be different interfaces," Aza said. The most obvious example takes us back to where I started with Tristan, in our very first conversation: there could be a button that says, "Here are all your friends who are nearby and are indicating they'd like to meet up today." You click it, you connect, you put down your phone and hang out with them. Instead of being a vacuum sucking up your attention and keeping it away from the outside world, social media would become a trampoline, sending you back into that world as efficiently as possible, matched with the people you want to see.

Similarly, when you set up (say) a Facebook account, it could ask you how much time you want to spend per day or per week on the site. You might name ten minutes, or two hours—it's up to you—and then the website could help you to achieve your goal. One way could be that when you hit that limit, the website could radically slow down. In tests, Amazon found that even one hundred milliseconds of delay in the pace at which a page loads results in a substantial drop-off in people sticking around to buy the product. Aza said: "It just gives your brain a chance to catch up to your impulse and [ask]—do I really want to be here? No."

In addition, Facebook could ask you at regular intervals: What changes do you want to make to your life? Maybe you want to exercise more, or take up gardening, or become vegetarian, or start a heavy-metal band. It could then match you up with other people nearby—friends, or friends of friends, or interested strangers in your

neighborhood—who say they also want to make that change and have indicated they are looking for the equivalent of gym buddies. Facebook would become, Aza says, "a way of socially surrounding yourself with the behavior that you want." A battery of scientific evidence shows that if you want to succeed in changing something, you should meet up with groups of people doing the same.

At the moment, they said, social media is designed to grab your *attention* and sell it to the highest bidder, but it could be designed to understand your *intentions* and to better help you achieve them. Tristan and Aza told me that it's just as easy to design and program this life-affirming Facebook as the life-draining Facebook we currently have. I think that most people, if you stopped them in the street and painted them a vision of these two Facebooks, would say they wanted the one that serves their intentions. So why isn't it happening? It comes back, Tristan and Aza said, to the business model. If right now these social-media companies made the changes that you just read about, they would lose an enormous amount of money. Within the existing economic structure of the companies, they can't do the right thing by your attention span or the wider society. This—above everything else—is the rock-solid reason why you have to change the business model, if you want to change the way social media affects us.

The business model can only be changed by regulation imposed on these companies by governments, they said. Then the changes I just described would cease to be impossible threats to the bottom line and start to become very exciting ways to tempt subscribers. At the moment there is a fundamental clash between your interests—to be able to focus, to have friends you see offline, to be able to discuss things calmly—and the interests of the social-media companies. With the introduction of a ban on surveillance capitalism and a move to a different business model, that clash ends. As Tristan put it, you'd be paying for the interests to be aligned between you and the product you use. Suddenly that team of Silicon Valley engineers

behind the screen wouldn't be working against you and your deeper intentions; they'd be working *for* you and trying to *serve* your deeper intentions.

One day, Aza said to me: "The fundamental thing is that no one likes the way that they are spending time or making decisions with the way technology currently is. It's hard to get from that hill to this hill, because we have to go through a valley. That's the role of regulation—to help making crossing that valley easier. But the hill on the other side is much, much nicer."

I found so much of what Aza and Tristan had taught me persuasive—but I was wary about their argument that we need to use the law to stop these companies from continuing as they are. I wondered if they were overstating the problem. When I spoke with Nir Eyal, he said: "Every generation has these moral panics, where we only want to look at the negative sides" of an issue. He told me, "Tristan is reading, literally verbatim, from the 1950s about the comic-book debate," when many people believed that children were being made violent by a new wave of gory comics. In the 1950s, "people like Tristan went to the Senate and told the senators that comic books are turning children into addicted, hijacked [zombies]—literally, it's the same stuff. . . . Today, we think of comic books as so innocuous."

On this basis, he argues—and here he's not alone—that the science that Tristan and Aza and other critics of the current tech business model draw on is incorrect. He believes that some of the social science I have drawn on in the past two chapters is garbled or wrong.

I'll give you one detailed example, so you get a sense of this controversy. Tristan argues YouTube is radicalizing people, based on an array of evidence I mentioned before. Nir responds by pointing to a recent study by the coder Mark Ledwich suggesting that, in fact, watching YouTube had a slightly *deradicalizing* effect on its users. Tristan, in response, directs people toward the Princeton academic

Professor Arvind Narayanan, and many other critics of this study, who say that the research Nir is citing here is worthless. Let's go through this, step by step. The people who say YouTube radicalizes you argue that this effect happens over time. You create a profile, you log in, and gradually YouTube builds up knowledge of your preferences, and to keep you watching, the content it feeds you gets more extreme. But the research Nir cites didn't study *any* logged-in users. All they did was go to a video on YouTube—say, Boris Johnson giving a speech—and without logging in, they looked at the recommendations that appeared along the side. If you use YouTube in this highly unusual way, the videos don't become more extreme over time, and it might be fair to say YouTube is deradicalizing. But huge numbers of YouTube users *do* log in. (We don't know exactly how many, because YouTube keeps that information secret.)

For every conceivable way the tech companies could be screwing with us, there is a back-and-forth like this, with Tristan and Nir each citing rigorous social scientists who have reached opposing conclusions. Tristan draws on academics from Yale and New York University and Harvard; Nir draws on academics like Professor Andrew Przybylski at Oxford University, who agrees with Nir that Tristan's warnings are overheated. So what's happening? It's not that either of them is being disingenuous—it's that measuring the changes these sites are triggering is really complicated, and hard to figure out. We have to be honest that we are making decisions based on a lot of uncertainty here. In the long sweep of history, there will likely be some areas where it turns out Nir is right, and some where Tristan is right. That still leaves us with a basic dilemma. Right now, we need to make choices about whether to let social-media companies continue behaving as they have been. We have to figure out the balance of risk.

There are two things that helped me make up my mind about what I think we should do next. One was a thought experiment, and the other was hard evidence from inside Facebook itself.

Let's imagine Nir is wrong, and we all follow his advice anyway—

we allow surveillance capitalism to continue getting us “fiendishly hooked,” with only light regulation. Then let’s imagine Tristan is wrong, and we all follow his advice anyway—we regulate the Big Tech companies to stop their invasive practices.

If Tristan is wrong and we still follow his advice, you would have been tricked into creating a world where you get targeted with a lot less advertising, you spend less, you get spied on less, and in return, you have to pay a small sum each month to subscribe to a few social-media companies, or those companies have in some way been taken over as public utilities run in our collective interests, like the sewers or the highways. Now imagine if we did what Nir wants. What happens if he’s wrong? What are we left with? Attention shrinks even more, political extremism expands, and the disturbing trends we see around us continue to rise.

The second thing that persuaded me was even more decisive. One day, in the spring of 2020, it was revealed what Facebook actually thinks about these questions, in private, when they think we will never be able to hear them. A large number of internal Facebook documents and communications were leaked to the *Wall Street Journal*. It turns out that behind closed doors, the company had responded to the claims that their algorithms had damaged our collective attention and helped the rise of Trump and Brexit by convening a team of some of their best scientists and tasking them with figuring out if this was really true, and if it was, to figure out what they could do about it. The unit was called Common Ground.

After studying all the hidden data—the stuff that Facebook doesn’t release to the public—the company’s scientists reached a definite conclusion. They wrote: “Our algorithms exploit the human brain’s attraction to divisiveness,” and “if left unchecked,” the site would continue to pump its users with “more and more divisive content in an effort to gain user attention and increase time on the platform.” A separate internal Facebook team, whose work also leaked to the *Journal*, had independently reached the same conclusions. They found that 64 percent of all the people joining extremist groups

were finding their way to them because Facebook’s algorithms were directly recommending them. This meant that across the world, people were seeing in their Facebook feeds racist, fascist, and even Nazi groups next to the words: “Groups You Should Join.” They warned that in Germany, one-third of all the political groups on the site were extremist. Facebook’s own team was blunt, concluding: “Our recommendation systems grow the problem.”

After carefully analyzing all the options, Facebook’s scientists concluded there was one solution: they said Facebook would have to abandon its current business model. Because their growth was so tied up with toxic outcomes, the company should abandon attempts at growth. The only way out was for the company to adopt a strategy that was “anti-growth”—deliberately shrink, and choose to be a less wealthy company that wasn’t wrecking the world.

Once Facebook was shown—in plain language, by their own people—what they were doing, how did the company’s executives respond? According to the *Journal*’s in-depth reporting, they mocked the research, calling it an “Eat Your Veggies” approach. They introduced some minor tweaks, but dismissed most of the recommendations. The Common Ground team was disbanded and has ceased to exist. The *Journal* reported dryly: “Zuckerberg also signaled he was losing interest in the effort to recalibrate the platform in the name of the social good . . . asking that they not bring him something like that again.” I read this and I thought of my friend Raul Santiago, in his favela in Rio, being terrorized by helicopters sent by the far-right government that was elected with the help of these algorithms—algorithms so powerful that Bolsonaro’s supporters responded to his victory by chanting, “Facebook! Facebook!”

I realized that if Facebook won’t stop promoting fascism—promoting Nazism in Germany—they will never care about protecting your focus and attention. These companies will never restrain themselves. The risks of letting them continue behaving the way they have are greater than the risks of overreacting. They have to be stopped. They have to be stopped by us.



I was daunted. For a while, I felt I had no idea how we could achieve such a goal. Many people go this far in the argument, and then sputter to a pessimistic halt. They say that, yes, this system is messing with us in terrible ways, but we'll just have to adjust, because nothing and nobody can stop it. We live in a culture where there is a sense of deep political fatalism at every turn. I saw this when I wrote my book about the war on drugs, *Chasing the Scream*, and I traveled all over the world talking about it. Especially in the U.S., I kept hearing: Yes, you're right that the drug war is a disaster and failure. (Over 80 percent of Americans agree.) Yes, you're right that decriminalization or legalization would be better. But no, it will never happen—so do you know a good lawyer or rehab facility for an addicted relative? Political pessimism keeps people trapped in a search for purely personal and individual solutions.

But here's the truth: this despair isn't just self-defeating; I think it's actually empirically wrong. I reminded myself—forces as powerful as the tech companies have been defeated many times in human history, and it always happens in the same way. It is when ordinary people form movements and demand something better, and they don't give up until they have achieved it. I know that could sound vague or idealistic, so I want to give a very practical example of a change that happened in my family, and very likely in your family, in the past three generations.

I am forty-one years old. My grandmothers were the age I am now in the year 1962. In that year, my Scottish grandmother, Amy McRae, was living in a working-class tenement in Scotland, and my Swiss grandmother, Lydia Hari, was living on a mountain in the Swiss Alps. Amy had been forced to leave school when she was thirteen, because nobody thought it was worth educating girls. While her brother stayed in education, she was sent to work cleaning toilets, which she did all her working life. She had wanted to work with homeless people, but in practice women were locked out of

jobs like that, and she was told to know her place as a woman and shut up. Lydia grew up in a Swiss village, and as a teenager she was constantly drawing and painting. She wanted to be an artist. She was told that girls couldn't be artists. She got married young, and was told to obey her husband. I would sit in their kitchen years later, when her husband would hold out an empty mug and yell "*Kaffee!*" (coffee), and she was expected to scramble to fetch it. She would sometimes sketch, but she said it made her depressed, because it reminded her of what her life might have been.

My grandmothers lived in societies in which women were excluded from almost all systems of power and almost all choices about their lives. In 1962 there were no women in the British cabinet, the U.S. cabinet, or the Swiss government. Women made up less than 4 percent of the members of the British Parliament and the U.S. Senate, and less than 1 percent of Switzerland's Federal Assembly, where women weren't even allowed to vote in seventeen out of the country's twenty cantons (including the one where my grandmother lived). This meant that the rules were written by men for men. American and British women were banned from getting mortgages or opening bank accounts unless they were married and had written permission from their husbands. Swiss women were banned from getting jobs at all without the written permission of their spouses. There were no domestic-violence shelters anywhere on earth, and it was legal everywhere for a man to rape his wife. (When, in the 1980s, there were moves to ban rape within marriage, one California Assembly member objected, saying, "But if you can't rape your wife, who can you rape?") In practice, men could beat their wives, because the police did not regard this as a crime, and they could molest their daughters, since it was so taboo to speak out about this that nobody ever went to the police to report it.

As I type out those facts, I keep thinking about my fifteen-year-old niece. Like her great-grandmother, she loves to draw and paint, and every time I see her doing it, I think of Lydia, doing the same thing in her Swiss village eighty-five years before. Lydia was told to

stop wasting her time and start serving men. My niece is told: You're going to be a great artist—let's start looking at art schools. My niece never met my grandmother, but I believe that Lydia would have been happy to know the ways in which feminism changed the world.

I know it's exceptionally irritating for a male to mansplain the topic in this way, especially when so much sexism and misogyny remains, and when women still face huge barriers. I know the advance of women's rights is far from won, and many of the advances that have been made are under threat. I know only one thing here that is definitely true: the difference between the lives of my grandmother and the life of my niece is a stunning achievement, and it happened for one reason, and one reason only. There was an organized movement of ordinary women who banded together and fought for it, and continued fighting even when it was really hard.

There are, of course, many differences between the fight for feminism and the fight for our focus. But nonetheless, I kept returning to this example in my mind for a very basic reason. The feminist movement teaches us that huge and seemingly immovable forces can be challenged by ordinary people—and that when they are, it can lead to real change. The concentrated power of men in 1962 was vastly greater than the power of Big Tech as I write this in 2021. Men controlled almost everything—every parliament, every corporation, every police force—and they had for as long as those institutions had existed. It would have been very easy, in that situation, to say: Nothing can change; give up; women will just have to learn to live a life of subordination. Many people are tempted to think that now, when they contemplate the huge forces stealing our focus. But that's the thing about the pessimistic belief that we are powerless and can't change anything. It's false.

Think about another historical example. I'm gay. In 1962, I would have been put in jail for it. Now I can get married. Homophobia ruled for 2,000 years, and then it didn't. The difference—the only difference—was a movement of ordinary people demanding an end to the forces thwarting their lives. I am free because the people who

came before us didn't give up; they got up. Again, of course, there are big differences between the struggle for equality for gay people, and this fight. But there is a key parallel: no source of power, no set of ideas, is so large it can't be challenged. Facebook would love us to believe that their power is impregnable and there's no point fighting for change because that never works. But these companies are as fragile as every other powerful force that was torn down in the end.

If we don't form a movement and fight, what is the alternative? Tristan and Aza warned me that right now, we are only at the start of what unregulated surveillance capitalism will do to us. It is only going to become more sophisticated and more invasive. They gave me lots of examples. Here's one. There's a technology that exists called "style transfer." If you use it, you can show a computer lots of paintings by Van Gogh and then you point it at a new scene, and it can re-create it in the style of Van Gogh. Aza told me how "style transfer" could quite soon be used against you or me: "Google today could read all of your Gmail, come up with a model that can mimic your style, and then sell that to an advertiser. [You, as the user] don't even know what's going on," but you will start to receive emails that are unusually welcoming and persuasive, because they sound just like you. Aza explained: "Even worse, they could look at all of your Gmail, look at all the emails you responded to quickly and positively, and learn that style. So [they] learn the style that is uniquely persuasive to you. There is nothing illegal about that. There are no laws to protect you against that. Is it breaking your privacy? They're not selling your data. They're just selling an asymmetric knowledge about how you work—even more than you know about yourself—to the highest bidder."

It's an asymmetry so extreme that it will hack vulnerabilities you don't even know are vulnerabilities. There are technological innovations coming that will make the current forms of surveillance capitalism look as crude as Space Invaders looks to a kid raised on Fortnite. Facebook, in 2015, filed a patent for technology that will be able to detect your emotions from the cameras on your laptop and phone.

If we don't regulate, Aza warns, "our supercomputers are going to test their way to finding all our vulnerabilities, without anyone ever stopping to ask—is that right? It'll feel to us a little bit like we're still making our own decisions," but it will be "a direct attack against agency and free will."

Tristan's mentor, Jaron Lanier—a veteran Silicon Valley engineer—told me he used to be a consultant for loads of dystopian Hollywood movies, like *Minority Report*, but he had to stop because he kept designing ever-more-frightening technologies to warn people of what was coming—and designers kept responding by saying: That's so cool; how do we make that?

"Sometimes I hear people say it's too late to make certain changes to the web or platforms or digital technology," James Williams told me. But the ax, he said, existed for 1.4 million years before anybody thought to put a handle on it. The web, by contrast, is "less than ten thousand days old."

We are, I realized, in a race. To one side there is the rapidly escalating power of invasive technologies, which are figuring out how we work and fracturing our attention. On the other side there needs to be a movement demanding technologies that work for us, not against us; technologies that feed our ability to focus, instead of fracturing it. At the moment, the movement for humane technology consists of a few brave people like Professor Shoshana Zuboff, Tristan, and Aza. They are the equivalent of the scattered bands of brave feminists of the early 1960s. We all need to decide—are we going to join them and put up a fight? Or are we going to let the invasive technologies win by default?