INTRODUCTION

A squirrel dying in front of your house may be more relevant to your interests right now than people dying in Africa.

—*Mark Zuckerberg*, Facebook founder

We shape our tools, and thereafter our tools shape us.

—*Marshall McLuhan*, media theorist

Few people noticed the post that appeared on Google's corporate blog on December 4, 2009. It didn't beg for attention—no sweeping pronouncements, no Silicon Valley hype, just a few paragraphs of text sandwiched between a weekly roundup of top search terms and an update about Google's finance software.

Not everyone missed it. Search engine blogger Danny Sullivan pores over the items on Google's blog looking for clues about where the monolith is headed next, and to him, the post was a big deal. In fact, he wrote later that day, it was "the biggest change that has ever happened in search engines." For Danny, the headline said it all: "Personalized search for everyone."

Starting that morning, Google would use fifty-seven *signals*—everything from where you were logging in from to what browser you were using to what you had searched for before—to make guesses about who you were and what kinds of sites you'd like. Even if you were logged out, it would customize its results, showing you the pages it predicted you were most likely to click on.

Most of us assume that when we google a term, we all see the same results—the ones that the company's famous Page Rank algorithm suggests are the most authoritative based on other pages' links. But since December 2009, this is no longer true. Now you get the result that Google's algorithm suggests is best for you in particular—and someone else may see something entirely different. In other words, there is no standard Google anymore.

It's not hard to see this difference in action. In the spring of 2010, while the remains of the Deepwater Horizon oil rig were spewing crude oil into the Gulf of Mexico, I asked two friends to search for the term "BP." They're pretty similar—educated white left-leaning women who live in the Northeast. But the results they saw were quite different. One of my friends saw investment information about BP. The other saw news. For one, the first page of results contained links about the oil spill; for the other, there was nothing about it except for a promotional ad from BP.

Even the number of results returned by Google differed—about 180 million results for one friend and 139 million for the other. If the results were that different for these two progressive East Coast women, imagine how different they would

be for my friends and, say, an elderly Republican in Texas (or, for that matter, a businessman in Japan).

With Google personalized for everyone, the query "stem cells" might produce diametrically opposed results for scientists who support stem cell research and activists who oppose it. "Proof of climate change" might turn up different results for an environmental activist and an oil company executive. In polls, a huge majority of us assume search engines are unbiased. But that may be just because they're increasingly biased to share our own views. More and more, your computer monitor is a kind of one-way mirror, reflecting your own interests while algorithmic observers watch what you click.

Google's announcement marked the turning point of an important but nearly invisible revolution in how we consume information. You could say that on December 4, 2009, the era of personalization began.

WHEN I WAS growing up in rural Maine in the 1990s, a new *Wired* arrived at our farmhouse every month, full of stories about AOL and Apple and how hackers and technologists were changing the world. To my preteen self, it seemed clear that the Internet was going to democratize the world, connecting us with better information and the power to act on it. The California futurists and techno-optimists in those pages spoke with a clear-eyed certainty: an inevitable, irresistible revolution was just around the corner, one that would flatten society, unseat the elites, and usher in a kind of freewheeling global utopia.

During college, I taught myself HTML and some rudimentary pieces of the languages PHP and SQL. I dabbled in building Web sites for friends and college projects. And when an e-mail referring people to a Web site I had started went viral after 9/11, I was suddenly put in touch with half a million people from 192 countries.

To a twenty-year-old, it was an extraordinary experience—in a matter of days, I had ended up at the center of a small movement. It was also overwhelming. So I joined forces with another small civic-minded startup from Berkeley called MoveOn.org. The cofounders, Wes Boyd and Joan Blades, had built a software company that brought the world the Flying Toasters screen saver. Our lead programmer was a twenty-something libertarian named Patrick Kane; his consulting service, We Also Walk Dogs, was named after a sci-fi story. Carrie Olson, a veteran of the Flying Toaster days, managed operations. We all worked out of our homes.

The work itself was mostly unglamorous—formatting and sending out e-mails, building Web pages. But it was exciting because we were sure the Internet had the potential to usher in a new era of transparency. The prospect that leaders could directly communicate, for free, with constituents could change everything. And the Internet gave constituents new power to aggregate their efforts and make their voices heard. When we looked at Washington, we saw a system clogged with gatekeepers and bureaucrats; the Internet had the potential to wash all of that away.

When I joined MoveOn in 2001, we had about five hundred thousand U.S. members. Today, there are 5 million

members—making it one of the largest advocacy groups in America, significantly larger than the NRA. Together, our members have given over $120 million in small donations to support causes we've identified together—health care for everyone, a green economy, and a flourishing democratic process, to name a few.

For a time, it seemed that the Internet was going to entirely redemocratize society. Bloggers and citizen journalists would single-handedly rebuild the public media. Politicians would be able to run only with a broad base of support from small, everyday donors. Local governments would become more transparent and accountable to their citizens. And yet the era of civic connection I dreamed about hasn't come. Democracy requires citizens to see things from one another's point of view, but instead we're more and more enclosed in our own bubbles. Democracy requires a reliance on shared facts; instead we're being offered parallel but separate universes.

My sense of unease crystallized when I noticed that my conservative friends had disappeared from my Facebook page. Politically, I lean to the left, but I like to hear what conservatives are thinking, and I've gone out of my way to befriend a few and add them as Facebook connections. I wanted to see what links they'd post, read their comments, and learn a bit from them.

But their links never turned up in my Top News feed. Facebook was apparently doing the math and noticing that I was still clicking my progressive friends' links more than my conservative friends'—and links to the latest Lady Gaga videos more than either. So no conservative links for me.

I started doing some research, trying to understand how Facebook was deciding what to show me and what to hide. As it turned out, Facebook wasn't alone.

WITH LITTLE NOTICE or fanfare, the digital world is fundamentally changing. What was once an anonymous medium where anyone could be anyone—where, in the words of the famous *New Yorker* cartoon, nobody knows you're a dog—is now a tool for soliciting and analyzing our personal data. According to one *Wall Street Journal* study, the top fifty Internet sites, from CNN to Yahoo to MSN, install an average of 64 data-laden cookies and personal tracking beacons each. Search for a word like "depression" on Dictionary.com, and the site installs up to 223 tracking cookies and beacons on your computer so that other Web sites can target you with antidepressants. Share an article about cooking on ABC News, and you may be chased around the Web by ads for Teflon-coated pots. Open—even for an instant—a page listing signs that your spouse may be cheating and prepare to be haunted with DNA paternity-test ads. The new Internet doesn't just know you're a dog; it knows your breed and wants to sell you a bowl of premium kibble.

The race to know as much as possible about you has become the central battle of the era for Internet giants like Google, Facebook, Apple, and Microsoft. As Chris Palmer of the Electronic Frontier Foundation explained to me, "You're getting a free service, and the cost is information about you. And Google and Facebook translate that pretty directly into money." While

Gmail and Facebook may be helpful, free tools, they are also extremely effective and voracious extraction engines into which we pour the most intimate details of our lives. Your smooth new iPhone knows exactly where you go, whom you call, what you read; with its built-in microphone, gyroscope, and GPS, it can tell whether you're walking or in a car or at a party.

While Google has (so far) promised to keep your personal data to itself, other popular Web sites and apps—from the airfare site Kayak.com to the sharing widget AddThis—make no such guarantees. Behind the pages you visit, a massive new market for information about what you do online is growing, driven by low-profile but highly profitable personal data companies like BlueKai and Acxiom. Acxiom alone has accumulated an average of 1,500 pieces of data on each person on its database—which includes 96 percent of Americans—along with data about everything from their credit scores to whether they've bought medication for incontinence. And using lightning-fast protocols, any Web site—not just the Googles and Facebooks of the world—can now participate in the fun. In the view of the "behavior market" vendors, every "click signal" you create is a commodity, and every move of your mouse can be auctioned off within microseconds to the highest commercial bidder.

As a business strategy, the Internet giants' formula is simple: The more personally relevant their information offerings are, the more ads they can sell, and the more likely you are to buy the products they're offering. And the formula works. Amazon sells billions of dollars in merchandise by predicting what each customer is interested in and putting it in the front of the virtual

store. Up to 60 percent of Netflix's rentals come from the personalized guesses it can make about each customer's movie preferences—and at this point, Netflix can predict how much you'll like a given movie within about half a star. Personalization is a core strategy for the top five sites on the Internet—Yahoo, Google, Facebook, YouTube, and Microsoft Live—as well as countless others.

In the next three to five years, Facebook COO Sheryl Sandberg told one group, the idea of a Web site that isn't customized to a particular user will seem quaint. Yahoo Vice President Tapan Bhat agrees: "The future of the web is about personalization . . . now the web is about 'me.' It's about weaving the web together in a way that is smart and personalized for the user." Google CEO Eric Schmidt enthuses that the "product I've always wanted to build" is Google code that will "guess what I'm trying to type." Google Instant, which guesses what you're searching for as you type and was rolled out in the fall of 2010, is just the start—Schmidt believes that what customers want is for Google to "tell them what they should be doing next."

It would be one thing if all this customization was just about targeted advertising. But personalization isn't just shaping what we buy. For a quickly rising percentage of us, personalized news feeds like Facebook are becoming a primary news source—36 percent of Americans under thirty get their news through social networking sites. And Facebook's popularity is skyrocketing worldwide, with nearly a million more people joining each day. As founder Mark Zuckerberg likes to brag, Facebook may be the biggest source of news in the world (at least for some definitions of "news").

And personalization is shaping how information flows far beyond Facebook, as Web sites from Yahoo News to the *New York Times*-funded startup News.me cater their headlines to our particular interests and desires. It's influencing what videos we watch on YouTube and a dozen smaller competitors, and what blog posts we see. It's affecting whose e-mails we get, which potential mates we run into on OkCupid, and which restaurants are recommended to us on Yelp—which means that personalization could easily have a hand not only in who goes on a date with whom but in where they go and what they talk about. The algorithms that orchestrate our ads are starting to orchestrate our lives.

The basic code at the heart of the new Internet is pretty simple. The new generation of Internet filters looks at the things you seem to like—the actual things you've done, or the things people like you like—and tries to extrapolate. They are prediction engines, constantly creating and refining a theory of who you are and what you'll do and want next. Together, these engines create a unique universe of information for each of us—what I've come to call a filter bubble—which fundamentally alters the way we encounter ideas and information.

Of course, to some extent we've always consumed media that appealed to our interests and avocations and ignored much of the rest. But the filter bubble introduces three dynamics we've never dealt with before.

First, you're alone in it. A cable channel that caters to a narrow interest (say, golf) has other viewers with whom you share a frame of reference. But you're the only person in your bubble. In an age when shared information is the bedrock of shared

experience, the filter bubble is a centrifugal force, pulling us apart.

Second, the filter bubble is invisible. Most viewers of conservative or liberal news sources know that they're going to a station curated to serve a particular political viewpoint. But Google's agenda is opaque. Google doesn't tell you who it thinks you are or why it's showing you the results you're seeing. You don't know if its assumptions about you are right or wrong—and you might not even know it's making assumptions about you in the first place. My friend who got more investment-oriented information about BP still has no idea why that was the case—she's not a stockbroker. Because you haven't chosen the criteria by which sites filter information in and out, it's easy to imagine that the information that comes through a filter bubble is unbiased, objective, true. But it's not. In fact, from within the bubble, it's nearly impossible to see how biased it is.

Finally, you don't choose to enter the bubble. When you turn on Fox News or read *The Nation*, you're making a decision about what kind of filter to use to make sense of the world. It's an active process, and like putting on a pair of tinted glasses, you can guess how the editors' leaning shapes your perception. You don't make the same kind of choice with personalized filters. They come to you—and because they drive up profits for the Web sites that use them, they'll become harder and harder to avoid.

OF COURSE, THERE'S a good reason why personalized filters have such a powerful allure. We are overwhelmed by a

torrent of information: 900,000 blog posts, 50 million tweets, more than 60 million Facebook status updates, and 210 billion e-mails are sent off into the electronic ether every day. Eric Schmidt likes to point out that if you recorded all human communication from the dawn of time to 2003, it'd take up about 5 billion gigabytes of storage space. Now we're creating that much data every two *days*.

Even the pros are struggling to keep up. The National Security Agency, which copies a lot of the Internet traffic that flows through AT&T's main hub in San Francisco, is building two new stadium-size complexes in the Southwest to process all that data. The biggest problem they face is a lack of power: There literally isn't enough electricity on the grid to support that much computing. The NSA is asking Congress for funds to build new power plants. By 2014, they anticipate dealing with so much data they've invented new units of measurement just to describe it.

Inevitably, this gives rise to what blogger and media analyst Steve Rubel calls the attention crash. As the cost of communicating over large distances and to large groups of people has plummeted, we're increasingly unable to attend to it all. Our focus flickers from text message to Web clip to e-mail. Scanning the ever-widening torrent for the precious bits that are actually important or even just relevant is itself a full-time job.

So when personalized filters offer a hand, we're inclined to take it. In theory, anyway, they can help us find the information we need to know and see and hear, the stuff that really matters among the cat pictures and Viagra ads and treadmill-dancing music videos. Netflix helps you find the right movie to watch in its vast catalog of 140,000 flicks. The Genius function of

iTunes calls new hits by your favorite band to your attention when they'd otherwise be lost.

Ultimately, the proponents of personalization offer a vision of a custom-tailored world, every facet of which fits us perfectly. It's a cozy place, populated by our favorite people and things and ideas. If we never want to hear about reality TV (or a more serious issue like gun violence) again, we don't have to—and if we want to hear about every movement of Reese Witherspoon, we can. If we never click on the articles about cooking, or gadgets, or the world outside our country's borders, they simply fade away. We're never bored. We're never annoyed. Our media is a perfect reflection of our interests and desires.

By definition, it's an appealing prospect—a return to a Ptolemaic universe in which the sun and everything else revolves around us. But it comes at a cost: Making everything more personal, we may lose some of the traits that made the Internet so appealing to begin with.

When I began the research that led to the writing of this book, personalization seemed like a subtle, even inconsequential shift. But when I considered what it might mean for a whole society to be adjusted in this way, it started to look more important. Though I follow tech developments pretty closely, I realized there was a lot I didn't know: How did personalization work? What was driving it? Where was it headed? And most important, what will it do to us? How will it change our lives?

In the process of trying to answer these questions, I've talked to sociologists and salespeople, software engineers and law professors. I interviewed one of the founders of OkCupid, an algorithmically driven dating Web site, and one of the chief

visionaries of the U.S. information warfare bureau. I learned more than I ever wanted to know about the mechanics of online ad sales and search engines. I argued with cyberskeptics and cybervisionaries (and a few people who were both).

Throughout my investigation, I was struck by the lengths one has to go to in order to fully see what personalization and filter bubbles do. When I interviewed Jonathan McPhie, Google's point man on search personalization, he suggested that it was nearly impossible to guess how the algorithms would shape the experience of any given user. There were simply too many variables and inputs to track. So while Google can look at overall clicks, it's much harder to say how it's working for any one person.

I was also struck by the degree to which personalization is already upon us—not only on Facebook and Google, but on almost every major site on the Web. "I don't think the genie goes back in the bottle," Danny Sullivan told me. Though concerns about personalized media have been raised for a decade— legal scholar Cass Sunstein wrote a smart and provocative book on the topic in 2000—the theory is now rapidly becoming practice: Personalization is already much more a part of our daily experience than many of us realize. We can now begin to see how the filter bubble is actually working, where it's falling short, and what that means for our daily lives and our society.

Every technology has an interface, Stanford law professor Ryan Calo told me, a place where you end and the technology begins. And when the technology's job is to show you the world, it ends up sitting between you and reality, like a camera lens. That's a powerful position, Calo says. "There are lots of

ways for it to skew your perception of the world." And that's precisely what the filter bubble does.

THE FILTER BUBBLE'S costs are both personal and cultural. There are direct consequences for those of us who use personalized filters (and soon enough, most of us will, whether we realize it or not). And there are societal consequences, which emerge when masses of people begin to live a filter-bubbled life.

One of the best ways to understand how filters shape our individual experience is to think in terms of our information diet. As sociologist danah boyd said in a speech at the 2009 Web 2.0 Expo:

> Our bodies are programmed to consume fat and sugars because they're rare in nature. . . . In the same way, we're biologically programmed to be attentive to things that stimulate: content that is gross, violent, or sexual and that gossip which is humiliating, embarrassing, or offensive. If we're not careful, we're going to develop the psychological equivalent of obesity. We'll find ourselves consuming content that is least beneficial for ourselves or society as a whole.

Just as the factory farming system that produces and delivers our food shapes what we eat, the dynamics of our media shape what information we consume. Now we're quickly shifting toward a regimen chock-full of personally relevant information. And while that can be helpful, too much of a good thing

can also cause real problems. Left to their own devices, personalization filters serve up a kind of invisible autopropaganda, indoctrinating us with our own ideas, amplifying our desire for things that are familiar and leaving us oblivious to the dangers lurking in the dark territory of the unknown.

In the filter bubble, there's less room for the chance encounters that bring insight and learning. Creativity is often sparked by the collision of ideas from different disciplines and cultures. Combine an understanding of cooking and physics and you get the nonstick pan and the induction stovetop. But if Amazon thinks I'm interested in cookbooks, it's not very likely to show me books about metallurgy. It's not just serendipity that's at risk. By definition, a world constructed from the familiar is a world in which there's nothing to learn. If personalization is too acute, it could prevent us from coming into contact with the mind-blowing, preconception-shattering experiences and ideas that change how we think about the world and ourselves.

And while the premise of personalization is that it provides you with a service, you're not the only person with a vested interest in your data. Researchers at the University of Minnesota recently discovered that women who are ovulating respond better to pitches for clingy clothes and suggested that marketers "strategically time" their online solicitations. With enough data, guessing this timing may be easier than you think.

At best, if a company knows which articles you read or what mood you're in, it can serve up ads related to your interests. But at worst, it can make decisions on that basis that negatively affect your life. After you visit a page about Third World

backpacking, an insurance company with access to your Web history might decide to increase your premium, law professor Jonathan Zittrain suggests. Parents who purchased EchoMetrix's Sentry software to track their kids online were outraged when they found that the company was then selling their kids' data to third-party marketing firms.

Personalization is based on a bargain. In exchange for the service of filtering, you hand large companies an enormous amount of data about your daily life—much of which you might not trust friends with. These companies are getting better at drawing on this data to make decisions every day. But the trust we place in them to handle it with care is not always warranted, and when decisions are made on the basis of this data that affect you negatively, they're usually not revealed.

Ultimately, the filter bubble can affect your ability to choose how you want to live. To be the author of your life, professor Yochai Benkler argues, you have to be aware of a diverse array of options and lifestyles. When you enter a filter bubble, you're letting the companies that construct it choose which options you're aware of. You may think you're the captain of your own destiny, but personalization can lead you down a road to a kind of informational determinism in which what you've clicked on in the past determines what you see next—a Web history you're doomed to repeat. You can get stuck in a static, ever-narrowing version of yourself—an endless you-loop.

And there are broader consequences. In *Bowling Alone*, his bestselling book on the decline of civic life in America, Robert Putnam looked at the problem of the major decrease in "social capital"—the bonds of trust and allegiance that encourage

people to do each other favors, work together to solve common problems, and collaborate. Putnam identified two kinds of social capital: There's the in-group-oriented "bonding" capital created when you attend a meeting of your college alumni, and then there's "bridging" capital, which is created at an event like a town meeting when people from lots of different backgrounds come together to meet each other. Bridging capital is potent: Build more of it, and you're more likely to be able to find that next job or an investor for your small business, because it allows you to tap into lots of different networks for help.

Everybody expected the Internet to be a huge source of bridging capital. Writing at the height of the dot-com bubble, Tom Friedman declared that the Internet would "make us all next door neighbors." In fact, this idea was the core of his thesis in *The Lexus and the Olive Tree*: "The Internet is going to be like a huge vise that takes the globalization system . . . and keeps tightening and tightening that system around everyone, in ways that will only make the world smaller and smaller and faster and faster with each passing day."

Friedman seemed to have in mind a kind of global village in which kids in Africa and executives in New York would build a community together. But that's not what's happening: Our virtual next-door neighbors look more and more like our real-world neighbors, and our real-world neighbors look more and more like us. We're getting a lot of bonding but very little bridging. And this is important because it's bridging that creates our sense of the "public"—the space where we address the problems that transcend our niches and narrow self-interests.

We are predisposed to respond to a pretty narrow set of

stimuli—if a piece of news is about sex, power, gossip, violence, celebrity, or humor, we are likely to read it first. This is the content that most easily makes it into the filter bubble. It's easy to push "Like" and increase the visibility of a friend's post about finishing a marathon or an instructional article about how to make onion soup. It's harder to push the "Like" button on an article titled, "Darfur sees bloodiest month in two years." In a personalized world, important but complex or unpleasant issues— the rising prison population, for example, or homelessness— are less likely to come to our attention at all.

As a consumer, it's hard to argue with blotting out the irrelevant and unlikable. But what is good for consumers is not necessarily good for citizens. What I seem to like may not be what I actually want, let alone what I need to know to be an informed member of my community or country. "It's a civic virtue to be exposed to things that appear to be outside your interest," technology journalist Clive Thompson told me. "In a complex world, almost everything affects you—that closes the loop on pecuniary self-interest." Cultural critic Lee Siegel puts it a different way: "Customers are always right, but people aren't."

THE STRUCTURE OF our media affects the character of our society. The printed word is conducive to democratic argument in a way that laboriously copied scrolls aren't. Television had a profound effect on political life in the twentieth century— from the Kennedy assassination to 9/11—and it's probably not a coincidence that a nation whose denizens spend thirty-six hours a week watching TV has less time for civic life.

The era of personalization is here, and it's upending many of our predictions about what the Internet would do. The creators of the Internet envisioned something bigger and more important than a global system for sharing pictures of pets. The manifesto that helped launch the Electronic Frontier Foundation in the early nineties championed a "civilization of Mind in cyberspace"—a kind of worldwide metabrain. But personalized filters sever the synapses in that brain. Without knowing it, we may be giving ourselves a kind of global lobotomy instead.

From megacities to nanotech, we're creating a global society whose complexity has passed the limits of individual comprehension. The problems we'll face in the next twenty years—energy shortages, terrorism, climate change, and disease—are enormous in scope. They're problems that we can only solve together.

Early Internet enthusiasts like Web creator Tim Berners-Lee hoped it would be a new platform for tackling those problems. I believe it still can be—and as you read on, I'll explain how. But first we need to pull back the curtain—to understand the forces that are taking the Internet in its current, personalized direction. We need to lay bare the bugs in the code—and the coders—that brought personalization to us.

If "code is law," as Larry Lessig famously declared, it's important to understand what the new lawmakers are trying to do. We need to understand what the programmers at Google and Facebook believe in. We need to understand the economic and social forces that are driving personalization, some of which are inevitable and some of which are not. And we need to understand what all this means for our politics, our culture, and our future.

Without sitting down next to a friend, it's hard to tell how the version of Google or Yahoo News that you're seeing differs from anyone else's. But because the filter bubble distorts our perception of what's important, true, and real, it's critically important to render it visible. That is what this book seeks to do.