the Facebook post. An individual should be allowed to decide whether she wants to be associated with this photograph, but the person uploading the photo should also make sure that those pictured want to appear online, tagged or not. These interdependencies can be difficult to negotiate and situational: a couple getting married may ask guests not to share updates about the event, for fear of offending those who aren't invited; a dinner party's host may say that he doesn't mind posts to Instagram, only for one of the guests to hide his face when the iPhone is pointed his way. Of course, with these same contextual barriers being eroded, and with practically any human feeling or endeavor now trackable and shareable, the sense of privacy as interdependent may be lost. Or it may be simply reconfigured: if anything is shareable, if we are intermittently transparent to one another but always transparent to ad networks and intelligence agencies, then perhaps our collective privacy is none at all.

FACEBOOK AND THE NEW NORM

No social-media firm has been as explicit about its desire to overturn popular notions of privacy as Facebook. In January 2010, during an on-stage interview at an industry awards event, Facebook's Mark Zuckerberg said that privacy could no longer be counted as a "social norm." Many companies would hesitate, he said, to institute privacy changes for 350 million users (the site's user base at the time). "We decided that these would be the social norms now," Zuckerberg crowed, "and we just went for it."

These "norms" have changed often. At a May 26, 2010 event, Zuckerberg promised to make privacy controls "simpler." The site's controls have hardly gotten easier to manage, and it's not uncommon for a user to institute a certain level of privacy only to return months later and find out that that option is no longer available. I once foraged through the hedgerow of Facebook's privacy controls and selected an

option so that none of my Facebook friends could see photos in which I was tagged. Now, not only can all of my friends see these photos, but I can also no longer find a setting within Facebook's privacy controls to hide my photos. But, as Joseph Turow notes, Facebook's privacy settings "are irrelevant when it comes to advertisers. In offering the data anonymously, Facebook claims the right to use even aspects of profiles that members have chosen not to make public."

Through these and other measures, Facebook's great achievement has been to repeatedly chip away at the edifice of privacy and ensure that each move—each removal of a privacy control, each introduction of a new feature that exposes more user information—is eventually accepted. We are all frogs in the Facebook pot, slowly being brought to a boil. In the view of writer and digital activist Cory Doctorow, "Facebook trains you to undervalue your privacy." As Doctorow indicates, this practice of undervaluing privacy is not so much a side effect as a core value. It's essential to Facebook's business model that its users feel less and less attachment to their privacy so that they can share more, churning out ever more data. Facebook's premium on frictionless sharing means that sharing should be natural and easy between users, but this lack of friction also applies to the process of Facebook's own information gathering.

The amount of data that Facebook collects on its users is enormous and would be the envy of any intelligence agency, if they didn't have access to it already. In 2011, Max Schrems, an Austrian law student, requested and received a copy of his data file from Facebook; it was 1,222 pages and contained information that Schrems hadn't intended to turn over to the social network, such as the geographic coordinates of where he logged in from, people he had unfriended, and other data he had deleted. Schrems went on to file a complaint with the office of the Irish Data Protection Commissioner (Facebook's European headquarters are in Ireland), claiming that they had violated various European privacy and data-protection laws. His action spurred the Irish government to audit Facebook's practices of data collection

and retention and recommend a number of changes. Facebook also claimed that it would introduce a site-wide policy of deleting some user data that was more than a year old. But absent vigorous activist campaigns such as Schrems's, such promises are rarely followed up on by independent auditing.

Facebook has made similar promises about the data it gathers through its ubiquitous Like buttons. The company's official justification for its use of the Like button as a tracking mechanism goes as follows: "We record some of this information for a limited amount of time to help show you a personalized experience on that site and to improve our products." The information collected in this manner is deleted or anonymized after ninety days. Facebook also says that this browsing information is not sold to third parties.

Do you trust them? Will it always be this way, or might Facebook, when its stock price starts flagging, decide to start retaining data longer or to sell it to some of the market-research firms that would love to get their hands on it?

Writer and technologist Anil Dash has written that the company is "advocating for a pretty radical social change to be inflicted on half a billion people without those people's engagement, and often, effectively, without their consent." (Dash was writing before Facebook membership surged past the 1 billion threshold.) These privacy policies, he warned, represent an ideology of radical transparency that can have unintended consequences from some users: "Facebook is philosophically run by people who are extremists about information sharing. Though I choose to talk about my politics, or my identity, or my medical history or my personal relationships, I can do so primarily because I have the privilege to do so thanks to my social standing, wealth, and the arbitrary fact of being born in the United States."

By putting small pieces of Facebook across the Web, the social network has essentially arrogated itself the right to watch and catalog us wherever we go. While we retain some ability to limit what we show to other people on Facebook, we have few ways to limit what Facebook

itself learns about us. And once that information ends up in the black box of Facebook's data centers, we have no idea how it might be used.

INTERNET TRACKING ENTERS THE PHYSICAL WORLD

In an age of fabulously cheap digital storage and data-as-a-commodity, there is little reason for social networks to stem their customer surveillance. Regulatory responses have been remarkably lenient: Google's $7 million fine for using its Street View cars to indiscriminately suck up financial and password information from unsecured home WiFi networks represented a rounding error for the company. (Previously, the FCC fined Google just $25,000 for obstructing its investigation into Street View.) And occasional flare-ups of user backlash have proved fleeting in the face of powerful, useful, and free services. With more than one billion active accounts, Facebook has built up a formidable network effect, in which the cost of opting out, for many users, is too high. Facebook and its peers have also seen little pushback from spreading their tracking mechanisms well beyond their own networks; they've become an integral part of the social web. In this way, Facebook can extend its logic of persistent user surveillance to the Internet writ large. A user may leave Facebook.com, but he remains under Facebook's careful watch. We become conditioned to Facebook's prying eyes, in the same way we became conditioned to Gmail reading our e-mail so as us to provide us contextual ads. (Imagine the backlash if the U.S. Postal Service started opening every letter, reading its contents, and inserting a contextually relevant advertisement. On the other hand, they do already scan the address information—the metadata—of every letter, cataloguing it for America's intelligence services.)

Some browsers, such as Google Chrome and Mozilla's Firefox, have installed Do Not Track features, which are supposed to stymie the ability of advertisers, targeters, and ad networks to track browsing

habits. But users must activate this capability in the browser's settings, and as of March 2013, only 11.4 percent of desktop Firefox users had activated Do Not Track. There's an even bigger flaw in this system, though: Web sites are under no obligation to respect these requests, and in fact, most don't. This Do Not Track capability was instituted after members of the advertising and tech industries, government officials, and privacy advocates came together to support the mechanism, which was a central part of the Consumer Privacy Bill of Rights that the Obama administration presented in February 2012. Nine months later, CNN said that "the entire plan is on life support"—a victim of faltering negotiations between privacy advocates, who claimed that Web giants such as Yahoo and AOL weren't negotiating in good faith, and the companies themselves, who said that the other side expected too much. In a meeting with the W3C, the international consortium that helps devise standards for the Web, a vice president of the Direct Marketing Association reportedly "proposed that Do Not Track signals should actually permit data collection for advertising purposes, the very thing the mechanisms were designed to control." The Association of National Advertisers then published an open letter to Microsoft CEO Steve Ballmer, criticizing his company for automatically enabling Do Not Track on its Internet Explorer 10 browser (which at the time hadn't even yet been released). Even the most cursory privacy measures, it seemed, would be vigorously contested.

But the effort was doomed from the beginning. The Privacy Bill of Rights called for corporations to sign up voluntarily, with enforcement entrusted to the congenitally toothless FTC. The bill itself was mostly a set of vague recommendations, along with some general principles, such as "Consumers have a right to secure and responsible handling of personal data." There was also some dispute over the meaning of Do Not Track, with *Wired*, for example, questioning whether logging which stories readers browsed in order to serve up recommendations counted as tracking.

In the absence of concerted industry action and meaningful government regulation, some people have taken measures into their own

hands. Coders have created a number of anti-tracking tools, often in the form of free browser plug-ins that users can install. Apps such as Ghostery, DoNotTrackMe, and Disconnect block the more than 2,000 "retargeters" that use cookies, ad networks, and other techniques to track your browsing history and present you related ads across the Internet. These plug-ins can also block surveillance from social widgets, including Like buttons. Some browser makers have also stepped up, with Safari and Firefox automatically blocking cookies from third-party sites that the user hasn't visited. But some of these blocking apps aren't quite what they claim to be. Evidon, the company that makes Ghostery, takes some of the data it collects from Ghostery users—there are eight million of them—and sells it to advertisers.

Anti-tracking and -targeting measures may provide short-term solutions for users seeking some modicum of privacy or anonymity while browsing the Internet, but they do little to overturn the industry's status quo, which remains single-mindedly focused on knowing more about user activities than the users themselves.

Critics claim that advertising is essential to the digital economy. It's what makes so many Web sites and services free. This may be true, but consumers have no responsibility to help support a broken, if widely used, business model. Whatever implied social contract existed between advertisers and users has been torn up by the industry. Never before has so much information been collected, so much commercial surveillance performed, on such a broad cross-section of consumers, with all of it digitized and freely traded among data brokers. As the current ardor for Big Data shows, information harvesting can be an essentially endless process, with the only limits being technological. As Helen Nissenbaum writes, "This faith in information, envisioned as an asset of enormous value, creates a virtually unquenchable thirst that can only be slaked by more information, fueling information-seeking behaviors of great ingenuity backed by determined and tenacious hoarding of its lodes. Inevitably, as our awareness of this landscape grows, so grows a sense of privacy under assault."

We console ourselves with bromides about how no one should expect privacy online—as if this is an unchangeable situation to which we should simply be resigned. At least, we're reminded, it's not like this out in the physical world. There, we are constantly bombarded with advertising that targets us but on a far more general, and less personalized, level—a Gucci ad in GQ caters to the magazine's readers' presumed interest in luxury goods; a billboard for the new "Iron Man" movie broadcasts more widely, hoping to interest any and all passersby (Hollywood blockbusters are mass products on the largest possible scale; nearly anyone is a potential customer). But that practice is changing. With privacy being increasingly leaky, contexts breaking down, offline and online networks intermingled, social-media accounts linked across devices and platforms, and advertising networks and companies such as Facebook buying up reams of consumer data, new possibilities in targeting are opening up, particularly with the addition of facial recognition software. We are moving toward a world in which the same kinds of technologies that track you online are now tracking your movements and behaviors in the physical world. And often, it's the same companies involved, insinuating their tracking and collection technologies into all aspects of your life.

TARGETING INDIVIDUALS

If you were walking down Oxford Street in west London in early 2012, you may have seen a bus-stop billboard featuring a 40-second video ad for a campaign to educate girls in the developing world. That is, if you were a woman. If you were a man (or recognized as one by the advertisement's built-in facial-recognition system), you would have instead been shown a shorter clip, encouraging you to visit the Web site of Plan UK, the organization behind the ad. The shorter ad was envisioned as a way of turning the tables on men, who usually have more choices and opportunities than women; only women, in this case, were allowed to

see the full message. In Japan, NEC has produced digital billboards that also recognize a subject's gender and market different products accordingly. An American company, Immersive, has done the same, touting its "software that turns any camera into an intelligent sensor." As these pieces of software have improved, they have also been able to gauge a customer's relative age and their reactions to the ad (how long they look at it, for example). The resulting analytics may be employed to further hone the campaigns. And the next step, of course, is to individualize the targeting process, so that the software will be able to match your face to a photo from one of your social-media profiles. The ad, then, wouldn't have to stop in the mall; it could follow you onto your cell phone or appear next time you log into Facebook. It might appear as you drive by a digital billboard and then continue the conversation on a TV screen in an office elevator. We would be told that these ads are simply the most relevant to us, that they are finely targeted to further our engagement or are responding to the interest we've shown in the product in the past. The ads that we're used to following us around the Internet would follow us throughout our world. This persistent targeting should be called what it really is: surveillance, stalking, harassment, visual pollution.

There are some new technologies that allow for limited targeting of public advertisements without violating user privacy or imparting a sense of being surveilled. Plan UK's ad comes close, but its use of facial recognition might be troubling to some. Consider another clever but more respectful use of targeting. In May 2013, the ANAR Foundation, a Spanish organization that aids children facing abuse, put up an advertisement featuring a large photo of a child's face and some words against child abuse. The advertisement, which was on street-level displays, such as bus stops, made use of Lenticular printing, which allows viewers to see different images from different angles. (Sometimes used for large movie posters, Lenticular printing can impart a sense of movement.) In this case, ANAR calculated the average height of a 10-year-old child and produced another image—with bruises on the

child's face and a message addressing children directly, along with a hotline number—that only people of that height could see. An adult, looking down from a higher angle, would see a basic message against child abuse. A small child, presuming he or she could read, would see a more targeted, urgent message. There are ways in which this technology might be misused or put toward more vulgar commercial ends, but in this case, it was an ingenious approach to spreading information in the public interest and to directing it toward those who need it most.

As wearable computing and the proliferation of digital displays, interfaces, sensors, and cameras bring down the wall between offline and online, the possibilities for tracking, targeting, and data collection increase immensely. Social networks and the logic they represent—persistent surveillance of users, industrial-level data collection—are becoming integrated into our surroundings. Some designers speak fancifully, but not improbably, about a future in which anything is potentially a display, where digital interfaces seamlessly appear and disappear, as needed, in the objects around us.

It's difficult to keep track of all the various programs and initiatives that tech companies have under way to monitor our activities. But here are some that will give you an understanding of the scope of the effort. We know that Facebook's Like buttons—similar to Twitter's social widgets—allow it to learn a great deal about your Internet activity, your life, your relationships, your personal history. The company has even worked on tracking where users move their cursors onscreen—for example, to see if they hover over certain ads but then decide not to click. Facebook's partnerships with the large data brokers Acxiom, Epsilon, and Datalogix allow it to know what you buy in retail stores, since these firms gather data about frequent-buyer cards, such as the ones you may use at CVS and the grocery store.

Google Now sifts through your smartphone data, calendar, e-mail, GPS, search, and many other sources of information to find out about your daily activities and keep you up-to-date with tips, directions, reminders, and advertisements. Google has tested using location information to detect when smartphone users enter retail stores in order to see if online searches (and ad impressions) lead to in-store visits. Depending on your device and software configuration—whether you've opted into your phone's location services, have Google apps on your iPhone, or (the easiest method) have an Android phone—Google may be able to monitor your location almost constantly. These sorts of efforts allow Google to tell advertisers that its mobile ads work and to serve you ads when they think you're most susceptible to them.

Stores, in turn, are working on tracking their customers like never before. Nordstrom, Bloomingdale's, American Apparel, Verizon, and other major retailers have utilized software that picks up on smartphones' Bluetooth and WiFi signals in order to monitor how customers move through stores. The data can be useful in assessing store design, how long customers spend in certain areas, the paths they take, and so on. Stores can also attach unique identifiers to each phone and track customers over repeated visits. The information can be sold to brands who want to know how consumers interact with displays. A clothing company might see that some female customers are repeatedly stopping in front of their new luxury line, lingering, and then leaving, leading them to believe that their prices are too high.

Other stores have used cameras and facial-recognition programs to gauge customers' moods and responses to different stimuli. An Italian company sells mannequins, called EyeSee, which contain cameras equipped with facial-recognition software that can recognize customers by approximate age, gender, and ethnicity. NEC has developed facial-recognition software that allows stores to recognize VIP customers when they enter. An app called Facedeals invites businesses to install facial-recognition cameras, which recognize customers based on their Facebook photos and then sends deals and coupons to their phones. The app, which also checks customers into the location, must be authorized by customers. The deals offered are based on the customers' "like histories," so "personalized deals can now be delivered to your smartphone from all participating locations—all you have to do

is show your face." SceneTap also places facial-recognition cameras in bars and uses them to track the ages and gender ratio of the patrons—data which is then viewable on maps in the SceneTap app and on its Web site. "These apps are bridgeheads, or perhaps trojan horses, for more powerful (and probably more intrusive) services to come," the technologist Alessandro Acquisti told *Ars Technica* after SceneTap's launch.

Google has received a patent for "pay-per-gaze advertising," physical advertisements with embedded sensors that can tell when customers are looking at them. It's perhaps the most literal example of the attention economy. Under such a system—which, Google's patent notes, can include "billboards, magazines, newspapers, and other forms of conventional print media"—advertisers wouldn't pay Google just based on number of impressions or clicks. They'd also pay Google every time you look at the ad. Your gaze becomes a metric of value, making it almost impossible to walk down the street and not be caught up in an economic exchange between two companies. Google's sensors are also supposed to be able to pick up on pupil dilation and other emotional cues, providing Google and its partners with information on how ads affect people on an instinctive level.

The same patent mentions Google's work on overlaying ads on Google Glass, perhaps the signal product reflecting Google's vision of providing each customer an intensely personalized experience. From search to ads to shopping to an automated digital personal assistant, Google is promising to shape your experience of the world around what it and its commercial partners believe you need to see.

Google has also introduced customized maps that are different for each individual user. Two designers who worked on the project claim that "the more context it has about you, the more useful it can be"—a constant refrain for supporters of these kinds of technologies. These maps are likely to be different each time they appear, meaning that no map is the same twice or the same for two different people. This might seem like an exciting use of user data, but it also raises some problems.

What happens when two people, perhaps two students working on an assignment, are separately looking at the same area, and Google decides to show them two different versions of it? Will my map be loaded with the places advertisers want me to see? Will Google decide that I don't need to see poor regions of my city, because they think I don't go there enough anyway, and my demographic data indicates I'm middle class? What other kinds of selection biases might occur? If Google thinks I haven't been to a park in a while—maybe it's winter or maybe I choose to leave my smartphone at home when I go out for a jog—will it stop showing me parks and other public spaces? Google already guides the routes I take when traveling. Perhaps, they'll decide to start directing me past restaurants that advertise with them and locations featuring billboards with their pay-per-gaze technology. As I pass these restaurants, I might receive ads or coupons in my Gmail inbox offering me a discount. Along the way, my entire urban experience potentially comes under the influence of Google.

The maps example also raises some of the same problems we run into when thinking about sorting algorithms. Knowing that each action can influence various overseeing algorithms, do we adjust our behavior accordingly? Do we do things just so that the algorithms monitoring us won't go off track, so that it'll still "like" the things we like? You might already do a form of this—say, give a thumbs-up to a song on Pandora because you want to hear more of that type (assuming you trust Pandora's system to usefully recognize a particular song type). You may start paying with cash at bars, worrying that, if your health insurer has access to credit card data, it might think those five beers were only for you, rather than you and your friends. You might rate art films, but not the new Marvel movie, so that you appear more cultured to the dating site offering you discounts and the Facebook app that lists what films you've recently seen.

From this point of view, these webs of tracking, advertising, and surveillance technologies are profoundly coercive—as surveillance tends to be. For better and worse, ads landscape our environments.

They influence the culture and invite us to view ourselves in certain ways. Someone who's constantly receiving advertisements for weight-loss drugs, gyms, and plus-size clothing will be receiving different messages than his neighbor who's shown ads for beach vacations, fancy watches, and the latest French bistro. With these systems working in concert, people living side by side, even in the same home and sharing the same devices, might be served far different information and guided to different stores and opportunities. (Some companies already claim to be able to recognize different users who spend time on the same gadgets.)

The tracking industry works with advertisers and social-media companies to leverage users' insecurities. In October 2013, a marketing firm called PHD released a report about when women feel least attractive, encouraging advertisers to target women with "quick beauty rescues" on Sundays and Mondays and social opportunities (including the expensive accessories and clothing society says they require) later in the week. The study even broke down by the hour when women are more likely to feel unattractive, with 5 a.m. to 7 a.m. being the peak, or trough, as the case may be. The ad network MediaBrix has developed a method for targeting gamers in what it calls "breakthrough moments" in mobile games. You might be struggling to crack a level on a strategy game on your phone, but when you finally do, an advertisement will pop up to congratulate you. MediaBrix also has products that introduce sponsored digital rewards (Congrats on beating that level! Here's a trophy from Gillette) or that allow gamers stuck at some point to choose to watch a video in order to overcome the obstacle. These methods are less instinctually repugnant than PHD's report on women's insecurities, but MediaBrix is operating on a similar principle: target users when they're at their most vulnerable.

Microsoft, too, has a patent for targeting users based on emotional states. The patent discusses examining what users are looking at, their perceived reactions, and serving up ads accordingly. Advertisers would also receive more control over their audience: "Advertisers

provide targeting data that includes the desired emotional states of users it intends to target," the patent reads. Samsung, whose Galaxy smartphones can already be controlled with eye movements, has been refining facial recognition technology. One of its patents would organize and regulate users' interactions based on their emotional states. The site *IPWatchdog* commented, "It appears that Samsung is seeking to improve social-media communications by limiting interactions between members who provoke negative emotions, or increasing interactions between members who instill positive emotions." Besides being able to push certain types of users into desired exchanges—which would be helpful for regulating unwanted speech and for keeping users happy and chained to the platform—this technology could be used to provoke certain emotional responses in users and target them accordingly.

These systems are far more manipulative than any market research done in the past because advertisers now have the ability to reach us at virtually any time. They also know far more about us than their predecessors ever did, while making us complicit in the process by encouraging check-ins, structuring data, location services, and other data production/sharing that is, we are told, designed to improve a service. A growing crop of biometric tools—sleep measurement apps, fitness monitors, the thumbprint reader introduced on Apple's iPhone 5S, the gene-sequencing service 23andme.com—means that corporations are set to know us at the physical, even genomic level. ("Your DNA will be your data," says one particularly creepy HSBC ad spotted at JFK airport.) They may even anticipate health problems before we realize we have them. Read your fitness tracker's terms of service agreement. Are they required to notify you if they detect a health problem? Do they reserve the right to sell your personal information to health insurers? An FTC study found that thirteen health and fitness apps sold data to seventy-six other companies, sometimes including users' names and e-mail addresses. Within weeks of being acquired by Facebook, Moves, a fitness app, announced that it would begin sharing data with its new

parent company. The change would allow Facebook, should it desire, to follow their users into the physical world, gaining information both about their movements and their physical health. It's easier to push ads at someone if you know when he's hungry, tired, or ill.

Although these new devices are in turn creating new methods of surveillance, the cookie remains a useful and widely used tracking technology. But soon it might be obsolete. Apple, Microsoft, Google, Verizon, and other companies are in an arms race to improve tracking technology. This is why Do Not Track buttons and industry-government negotiations over such rights are meaningless. Future tracking technologies will be integrated at the hardware level, making them harder to disable with software, while your face will serve as another kind of cookie, to be measured and parsed by CCTV and facial-recognition systems.

The MAC address—a unique identifier, similar to a serial number, stored in each cell phone's hardware—has already been deployed as a tracking mechanism. In August 2013, the city of London was forced to take twelve recycling bins off the streets after it was reported that the bins were tracking the movements of passersby by noting their MAC addresses. Renew, the company behind the bins, had been soliciting local businesses, offering them ad-targeting information about people walking through the area. On their busiest day, the bins identified 106,629 people, each of them, on average, more than eight times. The bins also had Internet-connected screens to show ads. *Quartz*, the business site that broke the story, noted that the company that developed the tracking "orbs" in the bins described its technology as "a cookie for the real world." While the bins were a commercial project, they were installed before the 2012 Olympics, when the UK government was instituting extraordinary security measures, including putting surface-to-air missiles on top of apartment buildings. If given access to these bins (or if access were obtained through other means), British intelligence services could have a real-time map of people, and their data-rich devices, walking through the area. Seen this way, these spying bins are

little different than the license-plate readers on police cars and buildings throughout the United States. They are designed to identify and track people moving through public space—a type of mass surveillance that should be considered anathema.

From locations to moods to the latest research on when women claim to feel insecure—all of this data is crunched in the service of parting individuals with their money, or in getting them to do the equivalent micro-labor, such as clicking on an ad, causing a tiny payment to pass between advertiser and ad network owner. Do this billions of times and you can be the next Google.

Any advantage helps, even small psychological cues. Sociologists and behavioral economists call these "nudges": subtle reminders or gestures that can help people make better decisions in their self-interest. This might be useful at times; perhaps you want your fitness monitor to light up with an alert or to contact your doctor if certain vital signs deviate from expected patterns. In the hands of marketers, nudges might tell your mapping app to steer you past a particular billboard or excitedly congratulate you when you finally, finally beat that level in Candy Crush. Or maybe credit card offers will start appearing when marketers think you're most impulsive or when they know that your checking account is looking low. As the process is perfected, the consumer would be relegated, Mark Andrejevic explains, "to the role of feedback mechanism in an accelerating cycle of production and consumption."

The underlying irony here is that consumers produce the information which, through this constant feedback system, helps steer their behaviors. This is done by our browsing, social networking, publishing posts, and other forms of data production, but also sometimes more deliberately—liking brand pages, structuring data, doing check-ins, requesting coupons and special offers. In perhaps the most direct example of this participation, Acxiom, one of the industry's largest data brokers, unveiled a Web site called AboutTheData.com, which was supposed to allow the company to get ahead of bad press generated by

industry practices and various governmental investigations. The site, which gives people the ability to opt out of tracking by Acxiom, also was intended to show that the company wasn't afraid to operate more transparently. But About The Data's most diabolically savvy feature is that consumers can "correct" errors in their profiles—in other words, they can improve and structure their personal data. This might, as Acxiom claims, give consumers a better experience, but it is also clearly to the company's advantage, allowing them to make consumers complicit in filling out their data profiles. And while the site lets users suppress some data, it also doesn't show them everything that the company has on them.

Joseph Turow, the author of *The Daily You: How the New Advertising Industry Is Defining Your Identity and Your Worth*, fears possibilities of social discrimination. In an interview, he explained that "it's simply the idea that increasingly companies will use data about us in order to make decisions about how important we are, and some people will win and some people will not." This might be, he added, how the world has always worked, but that doesn't mean that it *must* be that way, nor that we should delegate this power to machines. "What we have here is a winner/loser scenario that takes place algorithmically," he said, "basically through the tracking of people and the using of predictive methods to figure out who's important and who isn't, on definitions of people—they don't even know it's going on. But companies are defining us, constructing us and making decisions about our importance without even our having any clue that this is taking place in any serious way."

How might it look to be at the losing end of one of these decisions? It might not just be the depressing ads you receive or, compared to a less prosperous neighbor, being overcharged for items on a shopping site. You might be denied disability insurance because the insurer looked at your social-media profile and decided that you didn't look depressed. That's what happened to Nathalie Blanchard, a Quebec woman whose insurance benefits were revoked after she posted pho-

tos of herself at a birthday party and at the beach—excursions which her doctor recommended in order to help battle her depression. You could be denied a loan because a bank thinks that your small number of Facebook friends means that your life is unstable or that you are unreliable. That's how the financial services company Lenddo determines credit worthiness. High schools and universities—many already monitoring current and potential students, whether for purposes of discipline or admission—may decide to start using predictive analysis to determine which students may become violent. The city of Chicago used a similar system to make lists of people likely to commit or be victimized by violent crimes and then tasked police officers and social workers to target these individuals.

"People will worry how they relate to one another and to machines and may even change their behavior because they want to be treated better," Turow said, before adding that because we don't know the parameters of the algorithms judging us, we're "constantly guessing."

There's social damage done by these practices. They create a systematic disrespect for people's privacy. They privilege certain types of people over others. They make everyday people worry about what kind of information—including biometric data, that most personal and revealing kind of information—is being collected on them. And we may not even know when it's being used. For instance, Facebook's tag suggestions for photos draw on facial recognition. You may have disabled Facebook's facial recognition feature, but it could still have a faceprint of you that it could use toward its own ends. Google scans photos uploaded to Picasa and Google+ to detect child pornography and report violators to law enforcement agencies. That's an understandable use of this technology, but once private companies get in the business of seeking out criminal activity for law enforcement, it's worth asking how far these policies go. If a terrorist attack or some other crisis were to occur, would Google give government agencies broader access to its stores of user photos in order to help identify

purchase apps on Google Play." This ad campaign has come under some criticism. Microsoft has long been cast as a stodgy old man, if not an outright villain, of the tech scene. The company certainly engages in some data collection practices comparable to Google's—like most big tech companies, it's deeply reliant on the Web's surveillance infrastructure—and the subsequent revelations of its participation in the NSA's PRISM program have done nothing to burnish Redmond's image as a privacy defender. The ad campaign also has the unfortunate consequence of presenting privacy as a commodity, on the same commercial plane as, say, how much storage space the company's cloud storage service offers. Although appealing to market dynamics may be important for the privacy debate, a purely economic approach to privacy has tended to favor large corporate actors while leaving users and regulators sidelined. Should there be a populist backlash to social networks' privacy policies, it should be grounded, at least in part, on moral arguments—if only because, in an age of networked privacy and contextual collapse, privacy transcends traditional boundaries. It is not merely a problem of one company's practices, or even of the whole digital world, but of our entire surveillance-saturated society. If governmental surveillance can be opposed on legal *and* moral grounds, shouldn't Google or Facebook's surveillance of its users submit to the same reasoning? Were the argument left purely to the market, those with the deeper pockets are likely to win, as they have been for too many years now.

Big Data and the Informational Appetite

You are the sum total of your data.
No man escapes that.

—Don DeLillo, *White Noise*

To understand the depth of our privacy problem, we have to look at the ideological, economic, and cultural roles that data collection and data mining have assumed in recent years. There is now so much data produced on our behalves—about a terabyte per capita per year—that a major industry has arisen, one that fits familiar patterns of techno-utopian thinking. Big Data, as this emerging field is called, promises to take the incredible amount of data collected—browsing histories, sensor information from smartphones, GPS coordinates, social-media activity, purchasing information, medical reports—and turn it into useful insights. Big Data has found supporters in health care, insurance, scientific research, education, energy, and intelligence. While some commentators have argued that the utopian possibilities of Big Data are overblown, others offer more dire outlooks: "the surveillance possibilities of the technology," according to the director of the Human Dynamics Lab at MIT, "could leave George Orwell in the dust." At its most far-reaching, Big Data promises predictions about the behaviors of individuals and population groups, as well as to forecast anything