From Michael Lynch's The Internet of Us, 2016, New York: W.W. Norton

5

Who Wants to Know: Privacy and Autonomy

Life in the Panopticon

In 1890, Samuel Warren and Louis Brandeis published an article in the *Harvard Law Review* arguing for what they dubbed "the right to privacy." It made a splash, and is now one of the most widely cited legal articles in U.S. history. What is less known is what precipitated the article. The Kodak camera had just been invented, and it (and cameras like it) was being used to photograph celebrities in unflattering situations. Because of this newfangled invention, Warren and Brandeis worried that technology—and our unfettered use of it—was negatively affecting the individual's right to control access to private information. Technology seemed to be outstripping our sense of how to use it ethically.

They had no idea.

In the first part of this book, we've seen how some ancient philosophical challenges have become new again. We've grappled with whether "reasonableness" is reasonable and whether truth is a fantasy. But these old problems are only half the story. To really appreciate how we can know more but understand less, we need to recognize what is distinctive about how we know now. And a good place to start is with this simple fact: the things we carry allow us to know more than ever about the world, faster than ever. But they also allow the world to know more about us—and in ways never dreamed of by

Warren and Brandeis. Knowledge has become transparent. We look out the window of the Internet even as the Internet looks back in.

Most of the data being collected in the big data revolution is about us. "Cookies"—those insidious (and insidiously named) little Internet genies—have allowed websites to track our clicking for decades. Now much more sophisticated forms of data analysis allow the lords of big data, like Google and Amazon, to form detailed profiles of our preferences. That's what makes the now ubiquitous targeted ad possible. Searching for new shoes? Google knows—and will helpfully provide you with an ad showing a selection of the kind of shoes you are looking for the next time you visit nytimes.com. And you don't have to click to be tracked. The Internet of Things means that your smartphone is constantly spewing data that can be mined to find out how long you are in a store, which parts of the store you visit and for how long, and how much, on average, you spend and on what. Your new car's "black box" data recorder keeps track of how fast you are traveling, where you have traveled and whether you are wearing your seatbelt. That's on top of much older technologies that continue to see widespread use—such as the CCTV monitors that record events at millions of locations across the globe.

And, of course, data mining isn't done just for business purposes. Arguably, the United States' largest big data enterprise is run by the NSA, which was intercepting and storing an estimated 1.7 billion emails, phone calls and other types of communications *every single day* (and that was way back in 2010).[1] As I write this, the same organization is purported to be finishing the building of several huge research centers to store and analyze this

data around the country, including staggeringly large million-square-foot facilities in remote areas of the United States.

We all understand that there is more known about what each of us thinks, feels and values than ever before. It can be hard to shake the feeling that we are living in an updated version of Jeremy Bentham's famous panopticon —an eighteenth-century building design that the philosopher suggested for a prison. The basic idea was a prison as a fishbowl. Observation, Bentham suggested, affects behavior—and prisoners would control their behavior more if they knew their privacy was completely gone, if they could be seen by and see everyone at all times.

In some ways, our digital lives are fishbowls; but fishbowls we've gotten into willingly. One of the more fascinating facts about the amount of tracking going on in the United States is that hardly anyone seems to care. That might be due not to underreporting or lack of Internet savvy by the public (although both are true) but to the simple fact that the vast majority of people are simply used to it. Moreover, there are lots of positives. Targeted ads can be helpful, and smartphones have become effectively indispensable for many of us. And few would deny that increased security from terrorism is a good thing.

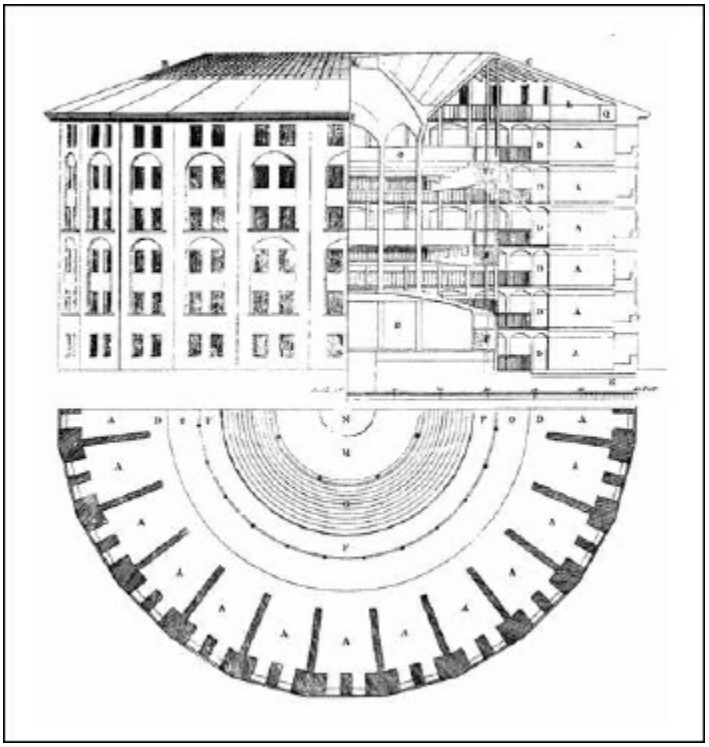

Fig. 2. Elevation, section and plan of Jeremy Bentham's Panopticon penitentiary, drawn by Willey Reveley, 1791.

Partly for these reasons, writers like Jeremy Rifkin have been saying that information privacy is a worn-out idea. In this view, the Internet of Things exposes the value of privacy for what it is: an idiosyncrasy of the industrial age.[2] So no wonder, the thought goes, we are willing to trade it away—not only for security, but for the increased freedom that comes with convenience.

This argument rings true because in some ways it *is* true: we do, as a matter of fact, have more freedom because of the Internet and its box of wonders. But, as with many arguments that support the status quo, one catches a whiff of desperate rationalization as well. In point of fact, there is a clear sense in which the increased transparency of our lives is not enhancing freedom but doing exactly the opposite—in ways that are often

invisible.

The Values of Privacy

If you are arrested for a serious crime in the United States today, your picture is taken, you are fingerprinted, and in some precincts the inside of your cheek is swabbed in order to obtain a sample of your DNA. In his dissenting opinion in the recent Supreme Court case on DNA identification techniques, Supreme Court Justice Antonin Scalia argued that such techniques amount to illegal searches.[3] We are, he said, opening our mouths to government invasion and tyranny.

Legally speaking, the case was complicated by a lack of clarity over whether collecting DNA constitutes a search. That is partly because DNA collection does not require a cheek swab; it can be collected from the skin or hair, for example. Thus, as the majority opinion noted, it is unclear why fingerprints wouldn't also constitute an illegal search if DNA samples do.

What *is* clear is that photos, fingerprints and DNA samples allow the police to identify and reidentify you—in ways that are increasingly immune to deception or alteration. You can change your name and your appearance, and fingerprints are not actually unique. But you can't as easily mess with the DNA—it is, in a real sense, part of who you are.

Of course, being able to identify criminal suspects is generally a very good thing—and DNA has proven to be an effective tool not only in this regard, but also for exonerating innocent people of crimes for which they've been falsely accused (and sometimes convicted). At the

same time, human beings have always been somewhat suspicious about new means of identification. That includes basic methods that we often overlook. Consider names. A rose may smell as sweet under any other name, but the fact that it has a name at all gives us an ability to reidentify it quickly and communicate that identification to one another. That's why, in some cultures, knowing someone's true name can give you (magical) power over them. You know how to identify "who they really are."

For similar reasons, images have often been said—as Warren and Brandeis were well aware—to have power. Even now, a photograph remains one of our best ways of identifying anything: we record in detail what our memories can't. Is it any wonder that individuals in some cultures were hesitant to allow their pictures to be taken? The idea that a camera could steal your spirit can be seen as a way of representing a real truth: that a picture identifies you, and like people's knowledge of your name is not something that is necessarily in your control.

This idea of control is closely connected to the idea of information privacy. The broad notion of privacy is difficult, if not impossible, to define in a straightforward way, and the narrower notion of information privacy is not much better. But even without a precise definition, it is clear that there are several marks or symptoms associated with information privacy. One of those concerns protection: we think of information as private to the extent that it is protected from interference or intrusion.[4] Another concerns control: information is private to the extent to which we control access to it.

Why do we value protecting and controlling our information? A cynic might say: we value it only when we have something to hide. But of course, even if this is true, it

doesn't really answer the question. That's because it depends on *what* you are hiding and *whom* you are hiding it from. Hiding a criminal past is one thing; hiding Jews in your basement from the Nazis is another.

In reality, there are much more basic reasons information privacy matters to us.

The Pool of Information

In the summer of 2014, following the revelations of Edward Snowden, the *Washington Post* revealed what many had long suspected: that the NSA, in targeting foreign nationals, is collecting and storing extremely large amounts of information on American citizens.[5] This information is not restricted to meta-data of the sort collected by the NSA's infamous phone data collection program. It is content—photos, Web chats, emails and the like.

U.S. law prevents the targeting of U.S. citizens without a warrant (even if it is just a warrant from the secret court established for this purpose by the Foreign Intelligence Surveillance Act (FISA) of 1978). But citizens' digital data is often vacuumed up "incidentally" when the NSA is collecting the posts, emails and so forth of legally designated foreign targets. Nothing currently prevents the NSA from engaging in this "incidental collection." And the incidentally collected data can be stored indefinitely. Moreover, no law prevents the agency—and other U.S. intelligence and law enforcement agencies—from accessing the incidentally collected content without a warrant, into perpetuity.

The storage of incidentally collected data seems

clearly wrong. Yet the reasons that make it so also help to explain why as a nation we sometimes sympathize with the sentiment voiced by Representative Mike Rogers in 2013: that your privacy can't be violated if you don't know about it—a non sequitur of such numbing grossness that only Peeping Toms could have greeted it with anything other than laughter.

But before getting back to the NSA, let's do another thought experiment. Imagine for a moment that I could perform something like the Vulcan mind meld with you (okay, okay, I'm dating myself). I telepathically read all your conscious and unconscious thoughts and feelings. You don't share your thoughts with me; I take them. I'm sure you'd agree that such an act of mental invasion would be wrong and harmful, but let's think about why.

The first and most obvious reason is that it has potentially dangerous consequences for you. And, Mike Rogers aside, that danger exists whether or not you know about my violation. Suppose you don't know. The more I know about you and the less you know about my knowledge, the easier it could be for me to take advantage of your ignorance: and the easier you will be to control or exploit.

Intentions matter. The fact that I can read your thoughts doesn't necessarily mean I will exploit you. I may be purely motivated by science: I might jot down your thoughts and do nothing to profit from them. Or I may be like Professor X of the comic book heroes X-Men, only motivated by truth, justice and the American way. More simply, the fact that I know what you like may help me guide you toward experiences you haven't had yet but would enjoy. Think of Amazon: part of their business model is predicated on acquiring information about their

customers' preferences—information often obtained without the customer really knowing it. They use this information to predict what else the customer might like—and to ensure that the customer is given every opportunity to buy it. There is no doubt that many of us feel uncomfortable about the amount of information collected now by corporations for the purpose of selling us stuff we need (or making us want to buy stuff others would like us to need). But it doesn't necessarily involve nefarious motives. So intentions matter, and even in the case of mind-reading, it is not absolutely certain that bad things will happen, even if you don't know about it.

But what if you *do* know that I am reading your thoughts? Well, you'll be wary, naturally. So wary that you will likely try and censor your thoughts and even your activities—perhaps by humming some Mozart in your mind to disguise your thoughts as best you can. And the reason you'd do so would be obvious—no matter how good I may seem to you now, you will want to minimize your exposure to exploitation and manipulation. This is not surprising. As Bentham knew when he designed his panopticon, observation affects behavior. But, of course, that too isn't necessarily bad. It is why security cameras are not always hidden. If you know you are being watched, you are less likely to act out. And that can be an instrument for good. Or not.

So, one reason privacy is important is that invasions of it can lead to exploitation, manipulation and loss of liberty. These, in turn, obviously can negatively affect a person's autonomy. But the possibility of bad effects is one thing, the actuality another. This is precisely the point that defenders of the NSA programs, for example, have been at pains to make. For all that's been said so far, there *might*

be negative consequences from, e.g., the NSA's policy of massive incidental collection and other data-sweep programs, *if* the agency or its architects were assumed to have bad or corrupt intentions. But why, some say, should we think that?

The fact is, however, that we don't have to know anything about the intentions of the program's architects in order to be worried. The NSA programs are dangerous to democracy even if we assume that their architects were motivated by the best of intentions—as no doubt many of them were. Roads to unpleasant places are frequently paved with the sweetest of intentions.

The NSA database could be described as a pool of information. This is an apt metaphor. In law, swimming pools are called attractive nuisances. They attract children and, as a result, if you own a pool, even if you are a watchful, responsible parent yourself, you still have to put up a fence. Similarly, even if we can trust that the architects of the NSA's various programs had no intention of abusing the information they are collecting about American citizens, the pool of information could easily prove irresistible. And the bigger the pool, the more irresistible it is likely to become. This is not just common sense, it explains why the NSA's repeated assertions that they aren't actually looking at the content of emails, or targeting Americans, should have been greeted with skepticism. The pool of data is a pool of knowledge. Knowledge is power; and power corrupts. It is difficult to avoid drawing the inference that *absolute knowledge might corrupt absolutely*.

That, not surprisingly, is the view of folks like Edward Snowden. But a growing number of stories strongly suggest that fear of abuse is more than a mere theoretical

worry. These examples are not constrained to the widely reported cases of NSA employees using their access to spy on sexual partners,[6] nor to similar cases in the UK where analysts collected sexually explicit photos of citizens without cause. More troubling, if less titillating, is the fact that the secret FISA court itself has complained that the NSA misrepresented its compliance with the court's previous rulings that various NSA techniques were unconstitutional.[7] In other words, the FISA court is being ignored by the very agency it is assigned to oversee and monitor. It is hard not to form the impression of an agency that feels it knows better than the judiciary or the Congress. And that, surely, should be worrying.

But the most disturbing fact is the massive continued storage of incidentally collected content itself (again: emails, photos, chat conversations and so on)—information that, as reported in the *Post*, is routinely searched by the CIA and the FBI—all without a warrant, even from the ineffective FISA court, and without any real oversight. Such searches needn't even be reported, and there is, presently, no legal oversight to prevent queries that are unrelated to national security, or even motivated by political ends. And relying on the agencies themselves to report abuses is like relying on the tobacco companies to tell us whether smoking is harmful.

While that's one of the major problems with the NSA collecting massive amounts of incidental information about Americans, it also helps explain why people don't seem too concerned. Putting up fences is arduous, time-consuming and expensive. And it does cut down on easy access to the water. So, if you want to get in that pool with the best intentions—you want to find the terrorists—it is

natural to think that the fence only gets in the way of what matters. If you trust that is what the owners of the pool are after, then worries about possible long-term negative consequences will seem abstract and, well, philosophical. After all, it is pretty clear human beings find it difficult to think about long-term consequences—that's true whether we are talking about swimming pools or global warming. If nothing bad has happened already that we know about as a result of privacy invasions, then what's the problem?

Unfortunately, if the pool of information about American citizens is systematically abused we aren't going to know about it—at least, not easily. When it comes to global warming, at least we'll get to realize the consequences of our current policies (or lack of them) one way or another. But the abuse of knowledge isn't going to be so obvious, and the abusers will have every reason to hide behind good intentions. That was one of the points made by the President's own review panel's report in 2013.[8] That panel—made up of not only writers and scholars including Cass Sunstein but former leaders of the CIA—suggested, in fact, more than simply fencing the pool (passing legislation to make it more difficult to access); they suggested the pool be drained. That is, they urged that all incidentally collected information (again, mostly on Americans, and far outweighing the amount being collected on warranted targets) simply be removed from the NSA's databases. This has not yet been done.[9]

Privacy and the Concept of a Person

The potential dangers of abusing big data are one reason the storage of incidentally collected information is

wrong. But there is another: the more insidious harm is not "instrumental" but "in principle."

Just this point was made over half a century ago in one of the most cited discussions of the right of privacy. In 1965, Edward J. Bloustein argued in a paper that what is wrong with such intrusions is

not the intentional infliction of mental distress but rather a blow to human dignity, an assault on human personality. Eavesdropping and wiretapping, entry into another's home, may be the occasion and cause of distress and embarrassment but it is not what makes these acts of intrusion wrongful. They are wrongful because they are demeaning of individuality and they are such whether or not they cause emotional trauma.[10]

Following what he took as the main point of Warren and Brandeis, Bloustein grounded the right of informational privacy on the intrinsic value of human individuality. The connection was what he called "personal freedom":

The fundamental fact is that our Western culture defines individuality as including the right to be free from certain types of intrusions. This measure of personal isolation and personal control over the doctrines of its abandonment is of the very essence of personal freedom and dignity, is part of what our culture means by these concepts.[11]

Let's unpack this thought. Philosophers have traditionally distinguished freedom of choice or action from what we might call the autonomy of decision. To see the difference, think about impulse buying. You may "freely" click on the Buy button in the heat of the moment —indeed, corporations count on it—without that decision

reflecting what really matters to you in the long run. Decisions like that might be "free" but they are not fully autonomous. Someone who makes a fully autonomous decision, in contrast, is committed to that decision; she owns it. Were she to reflect on the matter, she would endorse the decision as reflecting her deepest values.

Totally autonomous decisions are no doubt extremely rare; indeed, philosophers have long questioned whether they are possible at all. But it is clear that we value autonomy of decision, even if we can only approximate the ideal. That's because autonomy of decision is part of what it is to be a fully mature person. And that, I believe, tells us something about why privacy matters. It matters, at least in part, because information privacy is linked to autonomy, and thereby an important feature of personhood itself.

There are two ways to infringe on a person's autonomy of decision. The most obvious way is by *overruling* the decision, either by direct compulsion (I point a gun at your head) or by indirectly controlling your values and commitments (I brainwash you). A subtler way of infringing on your autonomy is to *undermine* it. Suppose a doctor makes the decision to give you a drug without asking your permission. Nobody has made you decide to do something. But your autonomy has been undermined nonetheless, and for an obvious reason: your decision has been foreclosed. You are not in a position to make the decision. It has been made for you.

Apply this to privacy. One mark of information privacy is control; that is, you have at least some *control over how and to whom you share those aspects of your self*. So consider a limit case. If you have a condition that compels you to say out loud every thought that comes into

your head—whether you like it or not—your autonomy of decision has been overruled. You are at the mercy of your condition.

But privacy invasions generally don't harm autonomy in this way. They don't overrule privacy. They undermine it. Suppose, to take a more old-fashioned example, that I break into your house and read your diary over and over again, every day. Suppose further that I make copies for my friends. Even if, again, you never learn of this, I am harming you in a new way—by undermining your capacity to control your private information. Whether you know it or not, that capacity is diminished. You may *think* you have the autonomy to decide whether to share your diary or not. But in fact, you are not in a *position* to make the decision; I've made that decision for you. Your autonomy of decision has been undermined.

As noted above, part of what makes your individual mind *your* mind is that you have a degree of privileged access to your mental states. And that includes, crucially, the ability to control access to the content of those thoughts and feelings—to choose whether and when you share this information with others. Part of the reason we value having this control is because it is a necessary condition for being in a position to make autonomous decisions, for our ability to determine who and what we are as persons.

It is here we see the danger inherent in systematic and sweeping collection by the government of the private information of citizens who have neither been charged with nor
suspected of a crime. Such intrusions on information privacy—
whether or not that information is acted upon or whether

the
intrusion is known—undermine not only dignity but, depending on how systematic the intrusions turn out to be, your actual capacity to control how and what information you share with others. The harm to your autonomy of decision becomes more global.

The systemic nature of the invasion therefore matters. Again, this point can be made starkly by looking back at our telepathic case, where I read your mind without your consent. Suppose it happens only once. Obviously, if I act on this knowledge in order to manipulate and control you, then I may directly harm your autonomy. And the risk of this happening may be great enough to prevent me from being tempted to use my power. Moreover, as I've been emphasizing, it shows a lack of respect for your status as a person, an autonomous being. But if I only read your mind—or your diary—once (and do nothing with the data), then presumably your autonomy itself is not affected. You don't become less in control of your self or face a diminished capacity of any sort.

Now let's suppose this happens again and again over time, in an organized way. From my perspective—the perspective of the knower—your existence as a distinct person will begin to shrink. Our relationship will be so lopsided that I may well cease to regard you as a full subject, as a master of your own destiny. As I learn what reactions you have to stimuli, why you do what you do, you may become like any other object to be manipulated *even if I do not, in fact, manipulate you.* You may be, as we say, dehumanized in my eyes. The connection between a loss of privacy and dehumanization is a well-known and ancient fact, and one which we don't need to appeal to

science fiction to illustrate. It is employed the world over in every prison and detention camp. It is at the root of interrogation techniques that begin by stripping people literally and figuratively of everything they own.

The connection between autonomy and privacy may sound surprising to some. After all, one could say: we are in fact willing to trade away our privacy as never before, precisely for the purpose of *increasing* autonomy. Our willingness—the thought goes—to trade privacy for security is just one example of this phenomenon. Another is our near total passivity when it comes to the trading of our data for profit by private corporations. We want more autonomy and they are providing it, by giving us convenience. Indeed, that's precisely the business model of corporations like Facebook or Amazon—to maximize convenience and anticipate our needs. Thus, one might say, it is not surprising that we click past all the privacy policies on the Web because we want the choices, the convenience—the autonomy—that only the playground of the infosphere can bring. Privacy suffers, but autonomy increases.

This argument, however, gets things the wrong way around. When we systematically collect private data about someone, we implicitly adopt what the philosopher Peter Strawson called the "objective" or detached attitude toward her.[12] We see her as something to be manipulated or controlled—even if, in fact, we never get around to the actual manipulating or controlling. Where privacy is limited in the detention camp or prison, the adoption of this attitude toward the inmate is of course explicit. It is an intrinsic feature of the enterprise and it is intuitively felt as such by those detained. Crucially, however, it remains implicit in more subtle invasions of privacy. In some cases,

this is unsurprising. When a business sells or otherwise profits from your private information—your Web searches, for example, or email address—it intentionally treats you as an object: an object of profit. Indeed, the nominal idea behind the privacy policies none of us read is to inform us of how our information will be used. They are a nod to our status as autonomous beings.

In truth, however, the Internet of Us is making privacy policies moot. When almost every object we interact with is wired, it becomes useless to assume that we consent to the mining of the data trail attached to our use of that object. That's because we simply have no way of being able to anticipate how the data being extracted from our refrigerators, for example, might be used in the future—by a company or by a government. Once the data is out there, it is out there. Any illusion we might have had about controlling or owning it gradually disappears. As Sue Halpern, an astute observer of the digital age, remarks: "The Internet of Things creates the perfect conditions to bolster and expand the surveillance state. In the world of the Internet of Things, your car, your heating system, your refrigerator, your fitness apps, your credit card, your television set, your window shades, your scale, your medications, your camera, your heart rate monitor, your electric toothbrush, and your washing machine—to say nothing of your phone—generate a continuous stream of data that resides largely out of reach of the individual but not of those willing to pay for it or in other ways commandeer it."[13]

Earlier I noted there are two marks to information privacy: control and protection. Control over our information may be increasingly under threat by the Internet of Things. But that only makes concentrating on

restricting and regulating information flow all the more important. The Internet of Things is enlarging the pool of data and information available for future use; that's why we need more fencing. We need the fences of regulations not only because they help prevent abuses, but because the pool threatens our autonomy.

There is another point here as well. Surveillance treats us as means, not as ends. And that is another reason the incidental collection of our data should worry us. A government that sees its citizens' private information as subject to tracking and collection has implicitly adopted a stance toward those citizens inconsistent with the respect due to their inherent dignity as autonomous individuals. It has begun to see them not as persons but as objects to be understood and controlled. That attitude is inconsistent with the demands of democracy itself.

Transparency and Power

Invasions of privacy aren't always wrong. If they were, we wouldn't have to spend so much time talking about the issue. My point is that they are always *pro tanto* wrong, as the legal scholars say. They are wrong—but wrong other things being equal.

Invasions of privacy can therefore be justified in the overall context. Searches of people's homes are judged "warranted" (that is, justified) for all sorts of reasons by the courts, as are surveillance operations of criminal suspects. Or consider the case of metal detectors and full body scanners at airports. The latter were (and still are) controversial on privacy grounds; moreover, more than one person argued that the scanner violated their dignity.

But while scans like this can make you uncomfortable, this sort of directed, publicly known invasion of one's privacy is not equivalent to the systematic program of incidental collection and meta-analysis of phone call data practiced by the NSA. That's because full body scans are given to commercial airplane passengers for a very specific reason: to detect whether they have a concealed weapon or explosives. This reason is well understood—or should be—by those given the scans. It is, in fact, a classic case of trading privacy for more security. It is a trade that may be justified, all things considered. Airport body scans are not stored indefinitely and open to the scrutiny of security agencies. They are made, examined, and eliminated. And they aren't being done secretly either. A better analogy would be this: secret scanners are set up so scans are taken of every person in his or her home. No one is told about the scans. They are stored indefinitely, and a wide range of agencies can examine them without a warrant. Still think that would be justified?

The possible negative consequences of losses of privacy in the digital age suggest that we must prepare for the worst even as we hope for the best. Think again of the swimming pool example. We need fences around our digital pools of information too. That's why, for example, some of the steps recently proposed by the Obama administration—to strengthen the FISA court's powers, and to limit some of the NSA's surveillance programs—are at least steps in the right direction.[14]

No one denies that governments naturally diminish our autonomy in all sorts of ways. Just participating in a government, as Hobbes stressed, is a trade-off. But the point I've been making in this chapter is that there is something different in the case of *systematic, unknown*

invasions of privacy. By invading our privacy without our knowledge, governments are making invisible decisions for the citizenry as a whole. That's not the same as restricting autonomy by asking people to go through a scanner at the airport. That's power visible to all, applied to all. Nor is it like wiretapping a particular citizen whom the courts have decided is a potential danger. Rather, these systematic, unknown invasions of privacy treat the citizenry as a whole in an unhealthy way. We are being regarded as unworthy of making up our own minds, whether we know it or not. That is an attitude that is corrosive of democracy, one made all the more corrosive by not being visible.

These reflections also give the lie to the idea that privacy of information is a modern creation. It is not. The source of privacy's value is deeper, lying at the intersection of autonomy and personhood itself. That is why privacy still matters. We are wise not to forget that, even as we trade it away.

Knowledge may be transparent, but power rarely is.