

The Architecture of Privacy

Lawrence Lessig[†]

Draft 2

Lessig 1998: This essay was presented at the Taiwan Net '98 conference, in Taipei, March, 1998.

[†] Jack N. and Lillian R. Berkman Professor for Entrepreneurial Legal Studies, Harvard Law School. Thanks to Tim Wu for extremely helpful comments on an earlier draft. Thanks to Professor Ching-Yi Liu for helpful comments on an earlier draft.

There is a part of anyone's life that is *monitored*, and there is a part that can be *searched*. The monitored is that part of one's day to day existence that others see, that others notice, that others could take note of, and that others could respond to, if response, in context, is appropriate. The searchable is the part of my life that leaves, or is, a record. As I walk down the street, my behavior is monitored. If I walked down the street in a small village in Mainland China, my behavior would be monitored quite extensively. This monitoring in both cases would be transitory — people would notice, for example, if I were walking with an elephant, or walking in a dress; but if there were nothing special about my walk, if I simply blend into the crowd, then I may be noticed for the moment, but would be forgotten soon after.

The searchable is less transitory. My scribblings in my diary leave a record of my thoughts. They can be searched. Stuff in my house is a record of what I possess. It too can be searched. And the recordings on my telephone answering machines are a record of who called, and what they said. It can be searched as well. These parts of my life don't so easily pass away. They are not in the same way ephemeral. They instead remain to be reviewed, if technology, and the law, permit.

This is an essay about privacy. My aim is to understand privacy through these two very different ideas. Privacy, in the sense that I mean here, can be described by these two different ideas. It stands in competition with these ideas. It is that part which is left after one subtracts, as it were, the monitored, and the searchable, from the balance of social life. Life where less is monitored is a life more private; and life where less can (legally perhaps) be searched is also a life more private. Thus understanding the technologies of these two different ideas — understanding, as it were, their architecture — is to understand something of the privacy that any particular context makes possible.

These architectures of privacy are many. There are many existing across the world today; there are many within any particular culture across its history. But I want to use this general way to describe architectures of privacy, as a general way to compare privacy across contexts. And in particular, as a way to see just why the context we are about to enter is so extraordinary different from any we have known.

For my claim today is that we are entering an age when privacy in any sense of that term will be fundamentally altered: An age when the extent of the monitored, and the reach of searchable, is far greater than any we have known. We can choose to let this change occur. Or we can choose to do something in response. After making plain the kind of change we should expect, my aim is to make understandable a range of responses, and to argue, if only implicitly, for some responses within that range.

The Monitored

So first let me describe a bit more completely this idea of the monitored. The monitored, as I described it, is that part of one's life that is watched. It is the part that is watched in an ordinary, or regular way. My focus here is not the infrequent spy, though if spying became extensive enough, spying would be part of the monitored. Nor is the periodic patrol of a cop part of the monitored. The monitored, as I mean it, is the regular, and persistent, watching of people or machines, whether the behavior watched is considered "public" or not.

Monitoring in social life is quite familiar. The most obvious historically has been the monitoring of a community. The image is commonplace: People living in a relatively small community, known by their neighbors, monitored as they come and go, as they buy in the market, as they associate at a local pub. Everything in that life, it is said, is known. Everything in that life, it is said, was known by others. One couldn't build in that world the modern liberal's conception of privacy. Privacy was what went on in one's head, not in one's life.

This is the picture that Americans often have of America at the founding. And it is the picture that leads many to say that there was no concept of "privacy" in America at the founding. Life then was life in public. One lived in small towns, everyone knew one's neighbors, everyone knew one's business. If you stayed out too late, or if you drank too much, or if you associated with the wrong people, or if you were rude to another in public — if you in any way breached an elaborate set of norms about how citizens were to behave, your breach would be noticed, and you would suffer the consequences of the breach. The social norms of such a society regulated individuals in that society, and they could therefore regulate much of the individual's life in such a society — since much of an individual's life was, in this sense, public, or in my terms, monitored.

This is a well-known account of life in early America, and a familiar account of life in much of America, and the world, today. The world is filled with places where individuals live life in just this sense — monitored by their neighbors, with this monitoring yielding social control. It is an important understanding of an important aspect of social life as well, for it emphasizes, properly I believe, the role that social norms can play in the regulation of social life. Social norms regulate; but they can regulate only if the behavior that they regulate can be monitored. This picture of life in a small town is a picture of life consistently monitored and this monitoring makes the control of social norms possible.

But this type of monitoring — the monitor of the small town, or the monitor of the community — has important features that we should not overlook. Features, that is, of its technology that mark it as a distinct kind of monitoring, importantly different from other more familiar instances of the very same idea.

The first feature is its relative transience; the second is who is doing the monitoring. My neighbors might remember that I was at the local market Saturday morning; they may even remember with whom I was talking; but they are not likely to remember exactly at what time, or everyone with whom I spoke. Nor will they know what I bought, or how much I paid, or whether I paid with large denomination bills or small. Of course, and again, if I did something out of the ordinary — if I brought my elephant to the market, or came with a woman who wasn't my wife — then my actions in this small town might be noticed in a less transitory way. Then my actions might be remembered. But in the ordinary case, they are not remembered. They are monitored for the moment, and then the record from that monitoring is forgotten.

More important than transience, however, is the check of who is doing the monitoring. Small towns, of course, have their busybodies — people who pry into the business of others — and they have their moral prudes — people whose standards are much stricter than most. But these enforcers of community norms are outliers. They define the extreme of a much narrower core. And it is this core of moral ideals that sets the limits on freedom that a community might define. The monitoring of a community serves this core; but its limit sets the limit on the burden of this monitoring. To an outsider, these norms might seem harsh. They might seem wrong. But for members of that community, they are just the sort of norms that the “ordinary” in that community obey. They are not extreme, or selective. They are not easily manipulated, or changed.

They are a set of influences that apply generally to like cases. And they get their force because they are applied by a community, acting as a normative community.

These are the features then of one particular architecture of monitoring. In a moment we will consider other architectures with different features. But before we consider these, consider the other side of this balance of privacy — the searchable. And consider it again, if you will, in the context of a small town, or, say, in early America.

As I've defined the term, the searchable is a function not only of what records there are that could be searched, but also of the technologies of searching, and the legal protections that might exist against the use of such technologies. Consider the technologies first.

In a word, they were crude. There was no simple way to hear, for example, a conversation going on between two people, locked securely in their own house. One might eavesdrop, but not easily and not with great success. Moreover, because life was on one's property, the law of trespass protected individuals from the wrongful entry of others — including, for example, the police. The searchable — letters, diaries, stuff in my house — was searchable only if the police got access to my property; the law protected me from their wrongful access, and the very nature of the architecture of property protected me from their wrongful access. Law, and the architecture of property, combined to establish a zone of privacy that neither the state, nor individuals, could easily breach.

The first American constitution guarantee that protection. The Fourth Amendment requires that searches be conducted only if reasonable; and that the warrant to search be granted only if there is probable cause to search. This constitutional affirmation of the value of privacy combined with legal protections that the common law long assured — protections again against trespass, or other invasions of privacy. These together gave legal support to the technological or architectural support for privacy that existed at the time.

One can understand, then, the scope of the searchable to be defined by two different factors. There is first the architecture of the social world that I am now describing — crude technologies for searching, relatively inefficient means of collecting data. These

inefficiencies themselves constituted a kind of a protection; they made it hard to search.

But as well as this architecture, the law was a protection. The law protected individuals against search; it limited the reasons the police could use for searching; it was a second line of defense against the invasion of prying eyes.

Privacy in this original context then was the product of this balance. On the one side, there is the life that was monitored by structures that support social norms. But on the other, there was the protection of law, and architecture, that combined to raise the costs of searching quite significantly. My life on the street might be monitored by my neighbors, but that monitoring produced few searchable records; and those records that were searchable were protected by both the architecture of property — that my walls were not made of glass, or I could lock my door — and by law, both constitutional and nonconstitutional. The balance of privacy then was this balance between the monitored, and the protections against search.

Preservation across contexts

As my story so far should make clear, much about this balance of privacy — at that time, and in anytime — depends upon the technology then existing. If what softens the burdens of monitoring is that monitoring is relatively transient, then technologies that eliminate transience increase the burden of monitoring. If what constitutes much of the protection of privacy in the home is that one who would breach has physically to enter the home, then technologies that allow invasion without physical invasion are technologies that reduce this privacy. Technologies in both cases can change; the question for law in both cases is how to respond to these changes. How to respond, that is, so that privacy is preserved.

In America, this question — how constitutional protections would respond to changing technologies — initially caused considerable struggle, and this struggle is quite relevant to the same questions today. As I said just minute ago, at the founding of my nation, the constitution protected the privacy of one's "person, papers and effects." That rule functioned like this: The common law protected individuals against trespass. If anyone — police officer or private citizen — entered my land without my permission, then I

had a right to a legal action against that person. If they entered my land without permission, I could sue them for trespass.

But sometimes the invader would have a defense. If the trespasser was a police officer, investigating a crime, then the officer would have a defense either if a jury determined that in the circumstances, the search was reasonable, or if the officer had a warrant to search. A warrant was permission granted by a judge to search, and it gave the police office immunity from the trespass law. So an officer deciding whether to search faced the following choice: If he searched without a warrant, then he could personally be liable for trespass if a judge did not believe the search reasonable; or he could secure from a judge a warrant, thereby insulating himself from liability. The regime thus created a strong incentive for the officer to secure a warrant, unless he was certain a search would be deemed reasonable.

The constitutional rule, therefore, takes this regime of trespass law for granted. It assumed that this law would continue, and it added just two qualifications on top of that law. The first was that regardless of the law, searches had to be reasonable. And second, it limited the circumstances under which a judge could grant a warrant. Only if the judge concluded that there was probable cause that a crime was committed could the judge issue a warrant.

So what happens, then, when searching can be accomplished without a trespass? For example, what happens when wire-tapping becomes possible? How would the Fourth Amendment protect privacy when privacy could be invaded without any trespass?

The question was first raised in the United States Supreme Court in 1928, in the case of *Olmstead v. United States*. In the midst of America's last great war on drugs — prohibition — the federal government start to use wire-tapping as a device for collecting evidence. State laws forbid wire-tapping, and the contracts that telephone companies had with their customers also promised that the wires would not be tapped. Nonetheless, the federal government began to tap phones, and in the case of *Olmstead*, the defendants challenged that wiretap on the grounds that it violated the 4th Amendment of the United States Constitution.

The Supreme Court disagreed. In an opinion by the Chief Justice, the court said that the 4th Amendment protected against trespass only; since wiretapping did not involve a trespass, the 4th Amendment did not protect against it. Hence evidence collected

through wiretapping would be admissible to convict Olmstead for violating the laws against prohibition.

Justice Brandeis, however, had a different view — a different view of the constitution, and a different view about the scope of the 4th Amendment. Certainly, Brandeis wrote, the constitution when originally authored protected only against trespass. But when it was authored, trespass was only an effective way to violate someone's privacy. When the constitution was enacted, to protect against trespass was to protect against the most significant invasions of one's privacy. But in 1928, that was no longer the case. In 1928, much of life had already moved onto the wires. And much of private life was now conducted on the telephone. In such a world, Brandeis argued, the protections of the 4th Amendment should be read to protect privacy on the phone as much as privacy in the home. To protect the same degree of privacy as the framers did, Brandeis argued, it was necessary to protect against more than trespass.

If there is a Justice who deserves the world's praise, if there is an opinion of the Supreme Court that should be the model for cyber-activists in the future, if there is a first chapter in the fight to protect privacy in cyberspace, it is this justice, and this opinion, and this case. Here, in as clear an example as any, is a method that will be central to cyberspace's survival as a place where values of individual liberty are sustained. Brandeis worked first to identify values from the original 4th Amendment, and then second, to *translate* these values into the context of cyberspace. Brandeis read beyond the specific applications the framers' had in mind, to find the meaning that they intended to constitutionalize; and in so reading, Brandeis found a way to read the constitution in the context of 1928 to preserve that meaning. Brandeis taught us to translate the framers' values into our interpretive context, in a way that had an extremely strong claim to constitutional fidelity.

However much one thought that Brandeis's method for reading the constitution was necessary in 1928 — however much one thought necessary a way of reading that took account of the changing contexts within which legal protections exist — it is impossible for us to avoid Brandeis' perspective today. We can't help but consider the technologies, or as I've called them, architectures of privacy in evaluating the world of privacy we are entering. For the world we are entering is about to change these architectures of privacy more completely and more extensively than any such change that we have seen to date.

We can begin to see this change by considering just a few stories. The first is a story about what is monitored; and the second, a story about the searchable.

Peter Lewis, writing in the *New York Times*, in an article titled “Forget Big Brother,” begins his story with the following account:

Surveillance cameras followed the attractive young blond woman through the lobby of the midtown Manhattan hotel, kept a glassy eye on her as she rode the elevator up to the 23d floor and peered discreetly down the hall as she knocked at the door to my room. I have not seen the videotapes, but I can imagine the digital readout superimposed on the scenes, noting the exact time of the encounter. That would come in handy if someone were to question later why this woman, who is not my wife, was visiting my hotel room during a recent business trip. The cameras later saw us heading off to dinner and to the theater — a middle aged married man from Texas with his arm around a pretty East Village woman young enough to be his daughter.¹

“As a matter of fact,” Lewis writes, “she is my daughter.”

Lewis’ is a story of the monitored — a hint to the emerging world of monitoring that is already constituting life in real space, and which promise even greater sway in cyberspace. Add to the cameras the credit card receipts, the telephone logs, the airplane tickets, the toll booths on the Triborough Bridge, the check in records at the hotel, the records from room service — add in all the records that get collected in the ordinary course in life in real space and the scope of real space monitoring begins to be clear.

Cyberspace of course will be even worse — or better, depending upon your perspective. Jerry Kang summarizes well the difference in an article soon to appear in the *Stanford Law Review*:

Imagine the following two visits to a mall, one in real space, the other in cyberspace. In real space, you drive to a mall, walk up and down its corridors, peer into numerous shops, and stroll through corridors of inviting stores. Along the way, you buy an ice-cream cone with cash. You walk into a bookstore and flip through a few magazines. Finally, you stop at a cloth-

¹ Peter H. Lewis, *Forget Big Brother*, *NYT* March, 19, 1998, pE1

ing store and buy a silk scarf, for a friend, with a credit card. In this narrative, numerous persons interact with you and collect information along the way. For instance, while walking through the mall, fellow visitors visually collect information about you if for no other reason than to avoid bumping into you. But such information is general (*e.g.*, does not pin-point geographical location and time of sighting), is not in a format that can be processed by a computer, is not indexed to your name or unique identifier, and is impermanent (in short-term human memory). You remain a barely-noticed stranger. One important exception exists: The scarf purchase generates data that are detailed, computer-processable, indexed by name, and potentially permanent.

By contrast, in cyberspace, the exception becomes the norm: Every interaction is like the credit card purchase. ... In this alternate universe, you are invisibly stamped with a bar code as soon as you venture outside your home. There are entities called "road" providers, who provide the streets and ground you walk on, who track precisely where, when, and how fast you traverse the lands, in order to charge you for your wear on the infrastructure. As soon as you enter the cyber-mall's domain, the mall begins to track you through invisible scanners focused on your bar code. It automatically records which stores you visit, which windows you peer into, in which order, and for how long. The specific stores collect even more detailed data when you enter their domain. For example, the cyber-bookstore notes which magazines you flipped through, recording which pages you have seen, for how long, and notes the pattern, if any, of your browsing. It notes that you picked up briefly a health magazine featuring an article on St. John's Wort, read for seven minutes a news weekly detailing a politician's sex scandal, and flipped every-so-quickly through a tabloid claiming that Elvis lives. Of course, whenever any item is actually purchased, the store as well as the credit, debit, or virtual cash company that provides payment through cyberspace take careful notes of what you bought--in this case, a silk scarf, red, expensive.²

In both accounts, the monitored increases. In both accounts, the scope of one's life subject to monitoring, changes. In both

² See Jerry Kang, *Cyberspace Privacy: A Proposal Regarding the Private Sector's Processing of Personal Information Generated in Cyberspace*, STANFORD LAW REVIEW (forthcoming 1998).

cases, this change is made possible by a change in the architecture of each space. The architecture is designed to capture data about our ongoing exchanges, and transactions, in ordinary life. Architectures, that is, designed to monitor.

The data that these systems collect, of course, is much like the data that a community in small town might collect. But again there are important differences. For unlike the data that the community might collect, the data from this monitoring is permanent, and searchable. It is not data collected and then discarded (since forgotten); it is data that is collected, and kept, and searchable by not just the community but by anyone who wants access to its facts. Think of the billions of gigabytes of email messages stored on across the world; or the tapes of telephone records archived by telephone companies; or the archives of frequent flyer miles, or credit card receipts, or calling card debits, or cash machines withdrawals, or toll booth records — think about all these, and you begin to get a sense of the extraordinary data that is coming to be collected as matter of routine, as a matter of what is ordinarily monitored.

This increase in the monitored is therefore increasing the searchable: More is monitored, this monitoring produces more that is searchable, and this more that is searchable remains waiting to be searched.

But as well as there being more to search, the costs of search are also falling. And, perhaps paradoxically, those falling costs reduce the legal protections against such search.

The first change is the more familiar, but we should divide those costs into two parts. One part are the costs borne by the searcher; the other the costs borne by the person being searched. The costs borne by the searcher are those costs involved in executing the search — the time spent searching, the expenses in executing the search, etc. The costs borne by the person being searched are not just the subjective costs of the search, but also the intrusion and disruption of the search.

Modern technologies are quickly reducing costs of both kinds. In real space, technologies such as telephoto lenses, or long distance microphones, or infrared cameras, or body scans, all make it cheaper to detect whatever the searcher is seeking. And in cyberspace, the change is all the more dramatic — as data moves onto a common protocol network, and systems for data matching become

all the more sophisticated. In both cases, the changes will mean a sharp reduction in the costs of a particular search, and hence an increase, in this aspect at least, of the searchable.

The same change is occurring with the costs borne by the person being searched as well. For these same devices — devices to scan bodies from a distance; devices to listen through walls from hundreds of feet away; searches of online data which the owner never notices; wiretapping — all of these, of course, while also being efficient techniques for the searcher, are also less burdensome for the person being searched.

But it is this second reduction that yields the paradox that I adverted to before. For this increase in the efficiency of a search not only reduce the economic cost on searching (and thereby increase the searchable); they are reduce the legal justifications for interfering with the searches.

The reasons are straightforward. In the ordinary case, the legal grounds for limiting the power of the state to search have been justifications grounded in the burdens that such searches impose on the person being searched. So that as these burdens on the person being searched are removed, there is less and less justification for limiting the state's right to search. Thus as the costs of searching fall, the legal grounds for restricting the search fall as well.

An example will make the point. I said earlier that the American constitution limits the right of the state to search. Searches, the constitution requires, must be "reasonable." So consider the following. Imagine a worm — a bit of computer code designed to works its way across the net and locate holes in the architecture of the net such that it can place itself onto the hard disks of computer users. The worm is designed not to do any damage. It doesn't attach itself to any system or application file. The worm instead simply places itself onto a hard disk, and searches that disk.

Say this worm were designed by the Federal Bureau of Investigation — America's federal police force. And say the worm was designed to search for a particular file — an illegal file, let's say, either a file with a national security document, or an illegal copy of some software code. The worm was designed to search disks without the user noticing; it did its work completely in background. If it found what it was looking for, it would report back to the FBI

the location of the file; if it didn't, it would simply destroy itself. The worm would not be able to search beyond its mandate.

Would such a worm violate the constitutional right of privacy? Well, I believe this is actually a very hard question. Certainly there is the sense of an invasion of property, but no longer is the 4th amendment tied to conceptions of property. The test under the 4th amendment now is simply whether the search is reasonable. Here, the search imposes no burden on the innocent; and only burdens the guilty. It is, in this sense, an efficient search. And because efficient, it is in effect increasing the range of the searchable. It is a general search, but because it imposes none of the costs of a general search, it might well be understood best as a reasonable search — like the sniff of a dog at the airport, except here there is not even the fear of the dog.

The worm is just an example. But it is an example that points to a more general point. More is being monitored; more can be searched cheaply; more can be searched without imposing any burden on the person being searched — searched efficiently, that is. Limits, in other words, on searching — both practical and legal — are being eroded. And the result of this erosion will be an ever increasing range of one's life that it is, at anytime in the future, the subject of discovery.

How should we understand this change? How should we understand its source? In the terms of the model I earlier described, we should its source first as flowing from a change in the architecture of the space. Its source, that is, is the change we will see in the architecture of a networked world. In real space, the default is that data is not collected. In real space, it takes effort — either the effort of a community, or the effort of a spy — to gather data. That is the architecture of the real world. And for most of our history, this architecture meant that any data so gathered was, in essence, useless. It was costly to hold; costly to use; costly to collect.

But the architecture of cyberspace is different. Or rather, the architecture is quickly becoming different. The architecture of cyberspace can be such that collecting data is the default. The world there can be made such that in the ordinary case, data is collected — invisibly, behind the scenes, efficiently, with no burden on the user. The data is collected; it is more easily searched; and the legal protections against its search — protections grounded in the burden that a search would create — disappear.

And so should we ask: Just how should we respond? How should we respond to this change in technology — to these changes in the architecture of cyberspace that yield a world unknown any we have known before.

The answer is not obvious, but if we put it in a regulatory context, some of the possibilities might become clear. That is my aim in the next section — to sketch a way of understanding this regulatory context, a model for understanding this problem of regulation. And in the final section, I'll use that model to help explain the differences in the responses of Europe and the United States, and to say something about the possibilities within each.

RESPONDING TO CHANGE

We should keep this issue in context. It is not as if the past two hundred years before the internet were years without technological change. It is not as if we have never faced this question before. Obviously, the question of individual privacy has been a dominant theme in legal thought for much of modern legal history. And plenty of nations have responded to the changes, by enacting legal proscriptions designed to replicate or create protections of an earlier period.

Some nations, but not my own. For while most modern democracies have enacted significant legal protections for privacy, my nation has not. Nations of Europe, and many democracies in Asia — including of course Taiwan — have enacted laws to create protections for privacy threatened by the emerging technologies of monitoring and search.

My nation, however, has been much slower to respond. My nation has been much more *laissez-faire* about privacy. We have no general federal statute protecting privacy, whether informational privacy or data privacy. We don't even have federal statutes effectively protecting medical privacy — the only group with that sort of protection of individuals in drug rehab clinics. Instead, where America has responded with law, America has responded with laws targeted in response to particular privacy problems. We have very effective protections for data about what videos people rent, for example, but only because a particular prominent American was embarrassed by the publication of the records of the videos he rented. American law is sporadic and partial — incomplete, from the perspective of data privacy in Europe, and inconsequential for most real protections.

The reasons for this lack of law in America protecting privacy are complex — one set relate to a general skepticism about law protecting this area; one set relates to the extraordinary lobbying power of interests that would use the data affected by informational privacy regulation; and one set relates to the claim that holders of this data respect the privacy of individuals in any case. But whatever the reason, we should not expect this feature of American law to change dramatically in the short term. Privacy in America is not about to be protected by law, in the way that privacy in Europe, and parts of Asia, is.

But does that mean that privacy won't be protected? Or put another way, is law the only kind of protection we might expect?

Think about the ways in which privacy is protected in real space — the many ways, and not just the protections of law. I want to focus just four. Law is one of those protections for privacy. The laws I have described, as well as state laws that supplement, as well as constitutional protections that supplement those as well. These combine, in the States as in Europe, to provide some protection for individual privacy — less here, I have argued than in Europe, but some nonetheless.

But laws are not the only protection for individual privacy. Norms protect privacy as well. At least among individuals, norms limit the kinds of questions one might ask, or the kinds of gossip one might listen to. And among corporations, norms restrict the kind of uses that these companies will make of the data they collect. These constraints are different from law — they get enforced, for example, not by the state, but by the sanctions of other members of a particular community. But they are nonetheless a source of constraint, functioning to protect privacy.

The market is a third type of constraint. Reputation in the market is affected by the use corporations make of privacy data, and in some cases, firms can offer more expensive services with a greater promise of privacy protection.

But in the story I've told so far, the most significant constraint protecting or possibly eroding, privacy, is the constraint of architecture. High walls make secure houses; sophisticated locks keep all but the most skilled burglar out; thick walls can't be listened through; thick curtains don't reveal. All these are features of the architecture of a particular space. And all these features in obvious ways increase, or extend, the privacy of a particular space, just as

other features of an architecture — a Panopticon, surveillance cameras, glass windows, offices without doors.

Now the point of describing these multiple constraints is to make obvious a perspective that is often lost in discussions of privacy, or of data privacy, today. That is to remember that it is these four constraints operating together that determine the privacy in any particular context. The four together could support privacy, or one could work against another. One might dominate the others; or two might work to the same end. To understand the privacy of a particular temporal context, or for a particular question, our focus must be on the mix of these four operating together.

It is against this background, then, that we should consider the state of data privacy in America today. For I've said already that laws in America are relatively slight, and are unlikely to be strengthened anytime soon. But given these alternatives to law, the real question should be whether these alternatives might supplement law to create a context in which privacy is protected.

One alternative, for example, would be norms. This is the solution of the Clinton Administration to the problem of data privacy. The administration wants industry to develop codes for regulating the handing of personal data. It wants industry to develop these codes on its own, and then enforce them without the involvement of the state. Industry would develop its own form of self regulation, and the state would rely on this self regulation to protect the privacy of its citizens.

There is much, of course, to be skeptical about with this solution — not the least of which being that the interests of commerce might well be different from the interests of the consumer. But it represents an alternative, the effectiveness of which must be considered when accounting for the interests protecting privacy.

A second alternative is architecture — technologies for recreating privacy where other technologies may have erased it. The most common example here is encryption — especially public key encryption, which would facilitate individuals hiding more effectively facts about themselves that they don't want third parties to know.

But encryption won't hide transactional data — it won't hide the monitoring of click streams, or telephone log records. And it won't easily hide the content of the records kept about us — except

to the extent those records are protected by those who collect that data. Moreover, encryption, oddly, increases the technologies of monitoring and search, for it facilitates an architecture within which identity can be established, and hence architectures which will require that identity be established.

Let me explain some more. Public Key encryption makes it easy to hide what one says. But it also makes it easy to authenticate who one is. Encryption facilitates both hiding, and authenticating, for the same technology that locks a conversation can be used to verify an identity. A digital signature, for example, can certify that I sent this, or a digital certificate certify that I am who I say I am. And it is this second part of the technology for encryption — this part that makes possible authentication — that we should consider when considering its effect on privacy.

As the cost of authenticating falls, we should expect the use of authenticating technologies to increase. As it is easier to say who I am, we should expect the growth of technologies that ask of me, who I am. The two will work together, for knowing who I am is valuable data. Thus it again will increase the data knowable, in a sense, by the system; it again is an architecture that will advance the ends of monitoring.

For this reason, I don't believe one can say — absolutely, or without qualification — that the development of encryption technologies will increase individual privacy. In the terms that began this essay, encryption may well reduce the searchable, by protecting what I hide; but by reducing the cost of authentication, it might well increase the monitored, and hence increase the searchable again. The technology, like many in this field, is janus faced — freedom enhancing from one perspective; control enhancing from another.

There are other technologies that might enhance privacy — other architectures that might make it possible to reclaim the protections of privacy. But before we consider these, consider first how the market might help.

My aim here is a sketch — there isn't the time here for me to do much more. My aim is to suggest how the market might be harnessed to this end of protecting privacy. Not on its own, but rather with the aid of a change in architecture. The mix, that is, of the market, and architecture, might well offer a solution to much of our present concern.

The intuition is this: Data is an asset. It is a resource which has become extremely valuable. And as it has become extremely valuable, commerce has tried to exploit it.

This use has a cost — an externality born by those who would rather this data not be used. So the trick is to construct a regime where those who would use the data internalize this cost. A regime to assure that they pay for this cost.

The laws of property are one such regime. If individuals can be given the rights to control their data, or more precisely, if those who would use data had first to secure the right to use it, then a negotiation could occur over whether, and how much, data should be used. The market, that is, could negotiate these rights, if a market in these rights could be constructed.

The advantages of the market are many, but the most important here is its ability to recognize diversity. A property regime gives the holder of the property right the power to hold out — until the buyer is willing to pay what the seller demands. But what this means is that people can hold out to different degrees. They can hold out for different amounts. What a property regime means is that people can sell for the price that suits them, regardless of what price might suit someone else.

The problem with this property regime, however, is its costs. The problem is the cost of negotiating the price to be paid. It would be impossible to imagine dickering with each click on the web. So how could this property regime be created?

It is here that the architecture comes into play; here, that is, that the change in the architecture I alluded to before comes into play. For there are a number of designs that code writers are proposing that might make this structures of negotiation possible.

One example is the regime of P3P, designed by the World Wide Web consortium. P3P is a standard for negotiating protocols on the web — a standard, that is, for negotiating protocols about privacy. It facilitates individuals setting the terms on which they will enter a site, for example, and then only entering sites that satisfy those terms. In the language of its authors, P3P is:

[an] interoperable way of expressing privacy practices and preferences by Web sites and users respectively. Sites' practices that fall within the range of a user's preference will be accessed "seamlessly," otherwise users

will be notified of a site's practices and have the opportunity to agree to those terms or other terms and continue browsing if they wish.³

The trick is a scheme that makes possible machine to machine communication. The web has made possible person to machine communication, and person to person communication. Architectures like P3P make possible machine to machine communication. And with machine to machine communication possible, machines can bear the cost of this negotiation. Machines, that is, could be our agents for protecting our privacy.

This solution again mixes both a market and architecture response. It is a solution that imagines the two working together, to create a kind of protection for privacy that law alone couldn't provide. If successful, it might well suffice in the protection of some individual data — not all, and certainly not for all purposes. But some, or perhaps enough, or certainly more than we now have.

CONCLUSIONS

I want to draw this together, and then to a close.

I've offered a way to reckon privacy in any given context. I've suggested that we think of it as a function of the monitored and the searchable. These I have said are increasing dramatically just now — much more is monitored, and much more is searchable. Both are increasing radically both in the context of real space and cyberspace. And we are fast entering an age where more can be known, and more efficiently collected, than at any time in our history ever.

These changes, I've said, are changes brought about by a change in architectures. Of the constraints that might protect privacy, I've argued that this constraint — architecture — has shifted most significantly. Its shift has an ambiguous quality — it makes possible an efficiency we have not before seen; and it makes likely an extent of monitoring we have not yet known.

My aim in this talk has been to offer a way to understand this change more generally, and to offer a way to understand how government might respond. One response of course is law — the response of the Europeans. But there are other responses beyond law

³ See <http://www.w3c.org/p3p>.

— the response of norms, and the market, and architecture. And my aim by sketching these is to suggest the complexity of response to these changes that the technologies make possible. There are protections beyond law, and cyberspace can help facilitate those protections.

That is the hopeful account. But I want now to end on note of skepticism, or better, anxiety, about where we are. For as much as we might envision a time when changes could restore a degree of privacy, we should not ignore the changes that are already occurring, and the vulnerability these changes will create just now.

For the lack of laws protecting data notwithstanding, governments are moving to take advantage of the efficiencies these new architectures facilitate. In Taiwan, for example, the government is developing smart card technologies, combining national insurance information, and identity information — including fingerprints — on a single card. These cards will also contain a digital signature, identifying the holder when used with a governmental data base. They are envisioned to be complete records for each individual — perfect identifications, and perfect links with that person's past. Efficient IDs — far better than the IDs we have today.

These efficiencies of course are valuable. But they beg for structures that check their use. They beg for structures built into the system, that might help assure that they don't become tools of misuse. As a balance to these advances, we must create structures that check these advances. Structures that build in checks on systems of control, to assure they control consistent with values of our tradition.

I am arguing that a kind of inefficiency should be built into these emerging technologies — an inefficiency that makes it harder for these technologies to be misused. And of course it is hard to argue that we ought to build in features of the architecture of cyberspace that will make it more difficult for government to do its work. It is hard to argue that less is more.

But though hard, this is not an argument unknown in the history of constitutional democracies. Indeed, it is the core of much of the design of many of the most successful constitutional democracies — that we build into such constitutions structures of restraint, that will check, and limit the efficiency of government, to protect against the tyranny of government. Edmund Burke, for example, said that the essence of a republic was that it would have

a senate to check both the excesses of democracy, and also the excesses of the executive. The senate was to be a balancing wheel in the structures of government. A governor on government; a structure designed to slow the effectiveness of both extremes within a democracy.

The same point helps explain much about the common constitutional rights in a constitutional democracy. They are, as John Perry Barlow has called them, “bugs” in the code of government: Elements designed to make government function less efficiently, so that rights are better protected. These “bugs” have value in contexts beyond the context of constitutional rights. They also have value in the very structure of government itself. One doesn’t want a perfectly efficient prosecutor, for fear that the prosecution will grow tyrannical. One doesn’t want an unimpeded executive, for fear that the executive will become arbitrary. One doesn’t want (France notwithstanding) a perfectly powerful and efficient legislature. One builds into a constitutional democracy limits on effectiveness of governmental power, to protect against abuse of governmental power.

Or at least one does so in some traditions. Or at least, this has been a concern of some traditions. And I understand well the claim of many that these traditions — these ideas of constitutionalism — are alien to others. I understand well the claim that in some traditions, this notion of individual rights is absent; this concept of individual over the community, misleading; this objective of preserving liberty over security, false. I understand this well not only because it is a familiar claim about cultures from this region, for example, but also because it has been a familiar claim in the American tradition. For again, remember, America didn’t begin its history as a libertarian state. The primary limits of the original constitution were limits on the *federal* government. States were not similarly limited. And local communities especially were not so limited — for these communities, as it is said today about much of life in eastern cultures, were communities where the primary regulator was a set of social norms that governed behavior. The primary regulator was not the state, but the community.

Thus, I understand the view that something greater than the individual might be considered greater than the individual. But I want to argue notwithstanding, that even if one believed that the community should remain the primary regulator — even if one rejected this notion of individual rights, or of the right to be different, or of Mill’s conception of individual liberty, of the duty of

general tolerance: Even if, that is, one reject the extreme libertarian view of individual liberty, and embraced instead a strong conception of community, or communitarianism — even then, one would not give have a reason to embrace the emerging technologies of control that I have described.

For the architectures of control that are emerging in this cyberworld are not the architecture of control of the traditional community. *Communities* are not, or would not be monitoring behavior, and enforcing norms through self-enforcement. This monitoring would be done by the state — by a small group separate from the community. And this separateness is extraordinarily significant, in two very different ways.

The first is about size. “The community” however one understood that term, is not the group that is controlling life in this emerging architecture of control. The group that gets the benefit of these architectures of control is the government. Governments, like guns, need not be bad; but when, like guns, they are placed in the wrong hands, they can become quite dangerous. And this is just what this power through knowledge means: that a small group has a great power, and that therefore, the risk of tyranny by this group is all the more great. The rules or requirements that can be enforced by this government are not necessarily the rules or requirements that would be enforced by the community. For they are not necessarily checked by the community. They instead get their power by pretending to enforce the will of the community, but instead get to enforce whatever will the small group might represent. They can stifle dissent — not because the community necessarily would, but because the architecture of control that has emerged gives them the power to monitor.

But a second difference is even more important. If we have learned anything about how communities function — if we have learned anything about the kinds of behavior that supports, or sustains a community, and the kinds of interventions that destroy it — then we have learned that for a community to sustain itself, it is the community itself that must enforce its rules. The norms of a community are sustained only so long as members of the community themselves are involved in the enforcement of those norms. Norms can’t be imposed externally, and in this contexts, governments are often external. If this enforcement is given to someone else — to the state, or to some other separate enforcing entity — then the community loses the practice of such enforcement. It loses the practice, and hence also loses the bonds that constitute

this community. In a weird but important paradox — only an inefficient community can sustain itself as a community; an efficient community (one that had institutions that efficiently enforced its norms) would self destruct. If members don't bear the cost of enforcing the rules of a community, the community will be lost.

What that means for the question of architecture is that the important question is not just what rules or norms are followed. More important is who actually enforces the norms that are to be followed. And so even if this extraordinary increase in the capacity to monitor means in the abstract that more norms could be enforced, unless they are enforced by the community, their enforcement will not increase strength of the community. They will be as external constraints — the constraints of an outsider, as the founders of my nation imagined the constraints of the federal government would be.

Thus even a “traditional” society — or better, a true traditional society, not one that uses tradition as a way to hide arbitrary power — has a reason to question these architectures of control. Or at least, has a reason to build into these architectures limits on the efficiency of such control. These limits are necessary not just to the ends of the individualistic society; they are important as well for avoiding tyranny in a society — an end shared by individualists, and communitarians alike.

My conclusion in the end is to urge us beyond a debate about privacy that is not really the appropriate debate in cyberspace. The issue in cyberspace is not the conflict between individualism and communitarianism; the issue in cyberspace is about whether local control of data in any sense will be permitted. Every community is vis-à-vis the world, an individual. And the question we must address about architectures is whether individuals of any kind will have control over the data the net makes available.

The net could flip the laws of nature as they apply to the collection of data. The “could” depends upon the architecture of the net. What is missing in discourse about cyberspace and its regulation is a richer understanding of the range of architectures that are possible. We must develop an attitude that thinks as critically about architectures as it thinks about laws; an attitude that understands the politics in both. We will only resolve finally and properly how this world should be made when we understand that we, in a critical sense, will be responsible for its making.

