

No doubt this mix will be controversial to some. But my aim is not so much to push any particular mix of settings on these modalities, as it is to demonstrate a certain approach. I don't insist on the particular solutions I propose, but I do insist that solutions in the context of cyberspace are the product of such a mix.

### Surveillance

The government surveils as much as it can in its fight against whatever its current fight is about. When that surveillance is human—wiretapping, or the like—then traditional legal limits ought to apply. Those limits impose costs (and thus, using the market, reduce the incidence to those most significant); they assure at least some review. And, perhaps most importantly, they build within law enforcement a norm respecting procedure.

When that surveillance is digital, however, then it is my view that a different set of restrictions should apply. The law should sanction "digital surveillance" if, *but only if*, a number of conditions apply:

1. The purpose of the search enabled in the algorithm is described.
2. The function of the algorithm is reviewed.
3. The purpose and the function match is certified.
4. No action—including a subsequent search—can be taken against any individual on the basis of the algorithm without judicial review.
5. With very limited exceptions, no action against any individual can be pursued for matters outside the purpose described. Thus, if you're looking for evidence of drug dealing, you can't use any evidence discovered for prosecuting credit card fraud.

That describes the legal restrictions applied against the government in order to enhance privacy. If these are satisfied, then in my view such digital surveillance should not conflict with the Fourth Amendment. In addition to these, there are privacy enhancing technologies (PETs) that should be broadly available to individuals as well. These technologies enable individuals to achieve anonymity in their transactions online. Many companies and activist groups help spread these technologies across the network.

Anonymity in this sense simply means non-traceability. Tools that enable this sort of non-traceability make it possible for an individual to send a message without the content of that message being traced to the sender. Implemented properly, there is absolutely no technical way to trace that message. That kind of anonymity is essential to certain kinds of communication.

It is my view that, at least so long as political repression remains a central feature of too many world governments, the government should recognize a protected legal right to these technologies. I acknowledge that view is controversial. A less extreme view would acknowledge the differences between the digital world and real world,<sup>39</sup> and guarantee a right to pseudonymous communication but not anonymous communication. In this sense, a pseudonymous transaction doesn't obviously or directly link to an individual without court intervention. But it contains an effective fingerprint that would allow the proper authority, under the proper circumstances, to trace the communication back to its originator.

In this regime, the important question is who is the authority, and what process is required to get access to the identification. In my view, the authority must be the government. The government must subject its demand for revealing the identity of an individual to judicial process. And the executive should never hold the technical capacity to make that link on its own.

Again, no one will like this balance. Friends of privacy will be furious with any endorsement of surveillance. But I share Judge Posner's view that a sophisticated surveillance technology might actually increase effective privacy, if it decreases the instances in which humans intrude on other humans. Likewise, friends of security will be appalled at the idea that anyone would endorse technologies of anonymity. "Do you know how hard it is to crack a drug lord's encrypted e-mail communication?" one asked me.

The answer is no, I don't have a real sense. But I care less about enabling the war on drugs than I do about enabling democracies to flourish. Technologies that enable the latter will enable the former. Or to be less cowardly, technologies that enable Aung San Suu Kyi to continue to push for democracy in Burma will enable Al Qaeda to continue to wage its terrorist war against the United States. I acknowledge that. I accept that might lead others to a less extreme position. But I would urge the compromise in favor of surveillance to go no further than protected pseudonymity.

### Control of Data

The problem of controlling the spread or misuse of data is more complex and ambiguous. There are uses of personal data that many would object to. But many is not all. There are some who are perfectly happy to reveal certain data to certain entities, and there are many more who would become happy if they could trust that their data was properly used.

Here again, the solution mixes modalities. But this time, we begin with the technology.<sup>40</sup>

As I described extensively in Chapter 4, there is an emerging push to build an Identity Layer onto the Internet. In my view, we should view this Identity Layer as a PET (private enhancing technology): It would enable individuals to more effectively control the data about them that they reveal. It would also enable individuals to have a trustable pseudonymous identity that websites and others should be happy to accept. Thus, with this technology, if a site needs to know I am over 18, or an American citizen, or authorized to access a university library, the technology can certify this data without revealing anything else. Of all the changes to information practices that we could imagine, this would be the most significant in reducing the extent of redundant or unnecessary data flowing in the ether of the network.

A second PET to enable greater control over the use of data would be a protocol called the Platform for Privacy Preferences (or P3P for short).<sup>41</sup> P3P would enable a *machine-readable* expression of the privacy preferences of an individual. It would enable an automatic way for an individual to recognize when a site does not comply with his privacy preferences. If you surf to a site that expresses its privacy policy using P3P, and its policy is inconsistent with your preferences, then depending upon the implementation, either the site or you are made aware of the problem created by this conflict. The technology thus could make clear a conflict in preferences. And recognizing that conflict is the first step to protecting preferences.

The critical part of this strategy is to make these choices machine-readable. If you Google "privacy policy," you'll get close to 2.5 billion hits on the Web. And if you click through to the vast majority of them (not that you could do that in this lifetime), you will find that they are among the most incomprehensible legal texts around (and that's saying a lot). These policies are the product of pre-Internet thinking about how to deal with a policy problem. The government was pushed to "solve" the problem of Internet privacy. Its solution was to require "privacy policies" be posted everywhere. But does anybody read these policies? And if they do, do they remember them from one site to another? Do you know the difference between Amazon's policies and Google's?

The mistake of the government was in not requiring that those policies also be understandable by a computer. Because if we had 2.5 billion sites with both a human readable and machine readable statement of privacy policies, then we would have the infrastructure necessary to encourage the development of this PET, P3P. But because the government could not think beyond its traditional manner of legislating—because it didn't think to require changes in code as well as legal texts—we don't have that infrastructure now. But, in my view, it is critical.

These technologies standing alone, however, do nothing to solve the problem of privacy on the Net. It is absolutely clear that to complement these technologies, we need legal regulation. But this regulation is of three very different sorts. The first kind is substantive—laws that set the boundaries of privacy protection. The second kind is procedural—laws that mandate fair procedures for dealing with privacy practices. And the third is enabling—laws that make enforceable agreements between individuals and corporations about how privacy is to be respected.

#### (1) Limits on Choice

One kind of legislation is designed to limit individual freedom. Just as labor law bans certain labor contracts, or consumer law forbids certain credit arrangements, this kind of privacy law would restrict the freedom of individuals to give up certain aspects of their privacy. The motivation for this limitation could either be substantive or procedural—substantive in that it reflects a substantive judgment about choices individuals should not make, or procedural in that it reflects the view that systematically, when faced with this choice, individuals will choose in ways that they regret. In either case, the role of this type of privacy regulation is to block transactions deemed to weaken privacy within a community.

#### (2) The Process to Protect Privacy

The most significant normative structure around privacy practices was framed more than thirty years ago by the HEW (Health, Education, Welfare) Advisory Committee on Automated Data Systems. This report set out five principles that were to define the "Code of Fair Information Practices."<sup>42</sup> These principles require:

1. There must be no personal data record-keeping systems whose very existence is secret.
2. There must be a way for a person to find out what information about the person is in a record and how it is used.
3. There must be a way for a person to prevent information about the person that was obtained for one purpose from being used or made available for other purposes without the person's consent.
4. There must be a way for a person to correct or amend a record of identifiable information about the person.

5. Any organization creating, maintaining, using, or disseminating records of identifiable personal data must assure the reliability of the data for their intended use and must take precautions to prevent misuses of the data.

These principles express important substantive values—for example, that data not be reused beyond an original consent, or that systems for gathering data be reliable—but they don't interfere with an individual's choice to release his or her own data for specified purposes. They are in this sense individual autonomy enhancing, and their spirit has guided the relatively thin and ad hoc range of privacy legislation that has been enacted both nationally and at the state level.<sup>43</sup>

### (3) Rules to Enable Choice About Privacy

The real challenge for privacy, however, is how to enable a meaningful choice in the digital age. And in this respect, the technique of the American government so far—namely, to require text-based privacy policy statements—is a perfect example of how not to act. Cluttering the web with incomprehensible words will not empower consumers to make useful choices as they surf the Web. If anything, it drives consumers away from even attempting to understand what rights they give away as they move from site to site.

P3P would help in this respect, but only if (1) there were a strong push to spread the technology across all areas of the web and (2) the representations made within the P3P infrastructure were enforceable. Both elements require legal action to be effected.

In the first edition of this book, I offered a strategy that would, in my view, achieve both (1) and (2): namely, by protecting personal data through a property right. As with copyright, a privacy property right would create strong incentives in those who want to use that property to secure the appropriate consent. That content could then be channeled (through legislation) through appropriate technologies. But without that consent, the user of the privacy property would be a privacy pirate. Indeed, many of the same tools that could protect copyright in this sense could also be used to protect privacy.

This solution also recognizes what I believe is an important feature of privacy—that people value privacy differently.<sup>44</sup> It also respects those different values. It may be extremely important to me not to have my telephone number easily available; you might not care at all. And as the law's presumptive preference is to use a legal device that gives individuals the freedom to be different—meaning the freedom to have and have respected wildly different subjective values—that suggests the device we use here is property. A property

system is designed precisely to permit differences in value to be respected by the law. If you won't sell your Chevy Nova for anything less than \$10,000, then the law will support you.

The opposite legal entitlement in the American legal tradition is called a "liability rule."<sup>45</sup> A liability rule also protects an entitlement, but its protection is less individual. If you have a resource protected by a liability rule, then I can take that resource so long as I pay a state-determined price. That price may be more or less than you value it at. But the point is, I have the right to take that resource, regardless.

An example from copyright law might make the point more clearly. A derivative right is the right to build upon a copyrighted work. A traditional example is a translation, or a movie based on a book. The law of copyright gives the copyright owner a property right over that derivative right. Thus, if you want to make a movie out of John Grisham's latest novel, you have to pay whatever Grisham says. If you don't, and you make the movie, you've violated Grisham's rights.

The same is not true with the derivative rights that composers have. If a songwriter authorizes someone to record his song, then anyone else has a right to record that song, so long as they follow certain procedures and pay a specified rate. Thus, while Grisham can choose to give only one filmmaker the right to make a film based on his novel, the Beatles must allow anyone to record a song a member of the Beatles composed, so long as that person pays. The derivative right for novels is thus protected by a property rule; the derivative right for recordings by a liability rule.

The law has all sorts of reasons for imposing a liability rule rather than a property rule. But the general principle is that we should use a property rule, at least where the "transaction costs" of negotiating are low, and where there is no contradicting public value.<sup>46</sup> And it is my view that, with a technology like P3P, we could lower transaction costs enough to make a property rule work. That property rule in turn would reinforce whatever diversity people had about views about their privacy—permitting some to choose to waive their rights and others to hold firm.

There was one more reason I pushed for a property right. In my view, the protection of privacy would be stronger if people conceived of the right as a property right. People need to take ownership of this right, and protect it, and propertizing is the traditional tool we use to identify and enable protection. If we could see one fraction of the passion defending privacy that we see defending copyright, we might make progress in protecting privacy.

But my proposal for a property right was resoundingly rejected by critics whose views I respect.<sup>47</sup> I don't agree with the core of these criticisms. For the

reasons powerfully marshaled by Neil Richards, I especially don't agree with the claim that there would be a First Amendment problem with propertizing privacy.<sup>48</sup> In any case, William McGeeveran suggested an alternative that reached essentially the same end that I sought, without raising any of the concerns that most animated the critics.<sup>49</sup>

The alternative simply specifies that a representation made by a website through the P3P protocol be considered a binding offer, which, if accepted by someone using the website, becomes an enforceable contract.<sup>50</sup> That rule, tied to a requirement that privacy policies be expressed in a machine-readable form such as P3P, would both (1) spread P3P and (2) make P3P assertions effectively law. This would still be weaker than a property rule, for reasons I will leave to the notes.<sup>51</sup> And it may well encourage the shrink-wrap culture, which raises its own problems. But for my purposes here, this solution is a useful compromise.

To illustrate again the dynamic of cyberlaw: We use law (a requirement of policies expressed in a certain way, and a contract presumption about those expressions) to encourage a certain kind of technology (P3P), so that that technology enables individuals to better achieve in cyberspace what they want. It is LAW helping CODE to perfect privacy POLICY.

This is not to say, of course, that we have no protections for privacy. As we have seen throughout, there are other laws besides federal, and other regulators besides the law. At times these other regulators may protect privacy better than law does, but where they don't, then in my view law is needed.

#### PRIVACY COMPARED

The reader who was dissatisfied with my argument in the last chapter is likely to begin asking pointed questions. "Didn't you reject in the last chapter the very regime you are endorsing here? Didn't you reject an architecture that would facilitate perfect sale of intellectual property? Isn't that what you've created here?"

The charge is accurate enough. I have endorsed an architecture here that is essentially the same architecture I questioned for intellectual property. Both are regimes for trading information; both make information "like" "real" property. But with copyright, I argued against a fully privatized property regime; with privacy, I am arguing in favor of it.

The difference is in the underlying values that inform, or that should inform, information in each context. In the context of intellectual property, our bias should be for freedom. Who knows what "information wants";<sup>52</sup> whatever it wants, we should read the bargain that the law strikes with holders

of intellectual property as narrowly as we can. We should take a grudging attitude to property rights in intellectual property; we should support them only as much as necessary to build and support information regimes.

But (at least some kinds of) information about individuals should be treated differently. You do not strike a deal with the law about personal or private information. The law does not offer you a monopoly right in exchange for your publication of these facts. That is what is distinct about privacy: Individuals should be able to control information about themselves. We should be eager to help them protect that information by giving them the structures and the rights to do so. We value, or want, our peace. And thus, a regime that allows us such peace by giving us control over private information is a regime consonant with public values. It is a regime that public authorities should support.

There is a second, perhaps more helpful, way of making the same point. Intellectual property, once created, is non-diminishable. The more people who use it, the more society benefits. The bias in intellectual property is thus, properly, towards sharing and freedom. Privacy, on the other hand, is diminishable. The more people who are given license to tread on a person's privacy, the less that privacy exists. In this way, privacy is more like real property than it is like intellectual property. No single person's trespass may destroy it, but each incremental trespass diminishes its value by some amount.

This conclusion is subject to important qualifications, only two of which I will describe here.

The first is that nothing in my regime would give individuals final or complete control over the kinds of data they can sell, or the kinds of privacy they can buy. The P3P regime would in principle enable upstream control of privacy rights as well as individual control. If we lived, for example, in a regime that identified individuals based on jurisdiction, then transactions with the P3P regime could be limited based on the rules for particular jurisdictions.

Second, there is no reason such a regime would have to protect all kinds of private data, and nothing in the scheme so far tells us what should and should not be considered "private" information. There may be facts about yourself that you are not permitted to hide; more important, there may be claims about yourself that you are not permitted to make ("I'm a lawyer," or, "Call me, I'm a doctor"). You should not be permitted to engage in fraud or to do harm to others. This limitation is an analog to fair use in intellectual property—a limit to the space that privacy may protect.