Lawrence Lessig, ***The Law of the Horse: What Cyberlaw Might Teach***
113 Harv. L. Rev. 501 (1999)

A few years ago, at a conference on the "Law of Cyberspace" held at the University of Chicago, Judge Frank Easterbrook told the assembled listeners, a room packed with "cyberlaw" devotees, * * * that there was no more a "law of cyberspace" than there was a "Law of the Horse"; that the effort to speak as if there were such a law would just muddle rather than clarify; and that legal academics * * * should just stand aside as judges and lawyers and technologists worked through the quotidian problems that this souped-up telephone would present. "Go home," in effect, was Judge Easterbrook's welcome * * *

Easterbrook's concern is a fair one. Courses in law school, Easterbrook argued, "should be limited to subjects that could illuminate the entire law. " "[T]he best way to learn the law applicable to specialized endeavors," he argued, "is to study general rules." This "the law of cyberspace," conceived of as torts in cyberspace, contracts in cyberspace, property in cyberspace, etc., was not.

My claim is to the contrary. I agree that our aim should be courses that "illuminate the entire law," but unlike Easterbrook, I believe that there is an important general point that comes from thinking in particular about how law and cyberspace connect.

This general point is about the limits on law as a regulator and about the techniques for escaping those limits. This escape, both in real space and in cyberspace, comes from recognizing the collection of tools that a society has at hand for affecting constraints upon behavior. Law in its traditional sense—an order backed by a threat directed at primary behavior—is just one of these tools. The general point is that law can affect these other tools—that they constrain behavior themselves, and can function as tools of the law. The choice among tools obviously depends upon their efficacy. But importantly, the choice will also raise a question about values. By working through these examples of law interacting with cyberspace, we will throw into relief a set of general questions about law's regulation outside of cyberspace.

I do not argue that any specialized area of law would produce the same insight. I am not defending the law of the horse. My claim is specific to cyberspace. We see something when we think about the regulation of cyberspace that other areas would not show us. * * *

Consider two cyber-spaces, and the problems that each creates for two different social goals. Both spaces have different problems of "information"—in the first, there is not enough; in the second, too much. Both problems come from a fact about *code*—about the software and hardware that make each cyber-space the way it is. * * * [T]he central regulatory challenge in the context of cyberspace is how to make sense of this effect of code. * * *

1. Zoning Speech.—Porn in real space is zoned from kids. Whether because of laws (banning the sale of porn to minors), or norms (telling us to shun those who do sell porn to minors), or the market (porn costs money), it is hard in real space for kids to

buy porn. In the main, not everywhere; hard, not impossible. But on balance the regulations of real space have an effect. That effect keeps kids from porn.

These real-space regulations depend upon certain features in the "design" of real space. It is hard in real space to hide that you are a kid. Age in real space is a self-authenticating fact. Sure--a kid may try to disguise that he is a kid; he may don a mustache or walk on stilts. But costumes are expensive, and not terribly effective. And it is hard to walk on stilts. Ordinarily a kid transmits that he is a kid; ordinarily, the seller of porn knows a kid is a kid, and so the seller of porn, either because of laws or norms, can at least identify underage customers. Self-authentication makes zoning in real space easy.

In cyberspace, age is not similarly self-authenticating. Even if the same laws and norms did apply in cyberspace, and even if the constraints of the market were the same (as they are not), any effort to zone porn in cyberspace would face a very difficult problem. Age is extremely hard to certify. To a website accepting traffic, all requests are equal. There is no simple way for a website to distinguish adults from kids, and, likewise, no easy way for an adult to establish that he is an adult. This *feature* of the space makes zoning speech there costly--so costly, the Supreme Court concluded in *Reno v. ACLU*, that the Constitution may prohibit it.

2. Protected Privacy.—If you walked into a store, and the guard at the store recorded your name; if cameras tracked your every step, noting what items you looked at and what items you ignored; if an employee followed you around, calculating the time you spent in any given aisle; if before you could purchase an item you selected, the cashier demanded that you reveal who you were--if any or all of these things happened in real space, you would notice. You would notice and could then make a choice about whether you wanted to shop in such a store. Perhaps the vain enjoy the attention; perhaps the thrifty are attracted by the resulting lower prices. They might have no problem with this data collection regime. But at least you would know. Whatever the reason, whatever the consequent choice, you would know enough in real space to know to make a choice.

In cyberspace, you would not. You would not notice such monitoring because such tracking in cyberspace is not similarly visible. * * * [W]hen you enter a store in cyberspace, the store can record who you are; click monitors (watching what you choose with your mouse) will track where you browse, how long you view a particular page; an "employee" (if only a bot) can follow you around, and when you make a purchase, it can record who you are and from where you came. All this happens in cyberspace--invisibly. Data is collected, but without your knowledge. Thus you cannot (at least not as easily) choose whether you will participate in or consent to this surveillance. In cyberspace, surveillance is not self-authenticating. Nothing reveals whether you are being watched, so there is no real basis upon which to consent.

These examples mirror each other, and present a common pattern. In each, some bit of data is missing, which means that in each, some end cannot be pursued. In the first case, that end is collective (zoning porn); in the second, it is individual (choosing privacy). But in both, it is a feature of cyberspace that interferes with the particular end. And hence in both, law faces a choice--whether to regulate to change this architectural

feature, or to leave cyberspace alone and disable this collective or individual goal. Should the law change in response to these differences? Or should the law try to change the features of cyberspace, to make them conform to the law? And if the latter, then what constraints should there be on the law's effort to change cyberspace's "nature"? What principles should govern the law's mucking about with this space? Or, again, how should law *regulate*? * * *

To many this question will seem very odd. Many believe that cyberspace simply cannot be regulated. * * * The anonymity and multi-jurisdictionality of cyberspace makes control by government in cyberspace impossible. The nature of the space makes behavior there *unregulable*.

This belief about cyberspace is wrong, but wrong in an interesting way. It assumes either that the nature of cyberspace is fixed--that its architecture, and the control it enables, cannot be changed--or that government cannot take steps to change this architecture.

Neither assumption is correct. Cyberspace has no nature; it has no particular architecture that cannot be changed. Its architecture is a function of its design * * * its code. This code can change, either because it evolves in a different way, or because government or business pushes it to evolve in a particular way. And while particular versions of cyberspace do resist effective regulation, it does not follow that every version of cyberspace does so as well. Or alternatively, there are versions of cyberspace where behavior can be regulated, and the government can take steps to increase this regulability.

To see just how, we should think more broadly about the question of regulation. What does it mean to say that someone is "regulated"? How is that regulation achieved? * * *

Behavior, we might say, is regulated by four kinds of constraints. Law is just one of those constraints. Law (in at least one of its aspects) orders people to behave in certain ways; it threatens punishment if they do not obey. The law tells me not to buy certain drugs, not to sell cigarettes without a license, and not to trade across international borders without first filing a customs form. It promises strict punishments if these orders are not followed. In this way, we say that law regulates.

But not only law regulates in this sense. Social norms do as well. Norms control where I can smoke; they affect how I behave with members of the opposite sex; they limit what I may wear; they influence whether I will pay my taxes. Like law, norms regulate by threatening punishment ex post. But unlike law, the punishments of norms are not centralized. Norms are enforced (if at all) by a community, not by a government. In this way, norms constrain, and therefore regulate.

Markets, too, regulate. They regulate by price. The price of gasoline limits the amount one drives--more so in Europe than in the United States. The price of subway tickets affects the use of public transportation--more so in Europe than in the United States. Of course the market is able to constrain in this manner only because of other constraints of law and social norms: property and contract law govern markets; markets operate within the domain permitted by social norms. But given these norms, and given

this law, the market presents another set of constraints on individual and collective behavior.

And finally, there is a fourth feature of real space that regulates behavior—"architecture." By "architecture" I mean the physical world as we find it, even if "*as we find it*" is simply *how it has already been made*. That a highway divides two neighborhoods limits the extent to which the neighborhoods integrate. That a town has a square, easily accessible with a diversity of shops, increases the integration of residents in that town. That Paris has large boulevards limits the ability of revolutionaries to protest. That the Constitutional Court in Germany is in Karlsruhe, while the capital is in Berlin, limits the influence of one branch of government over the other. These constraints function in a way that shapes behavior. In this way, they too regulate.

These four modalities regulate together. The "net regulation" of any particular policy is the sum of the regulatory effects of the four modalities together. A policy trades off among these four regulatory tools. It selects its tool depending upon what works best.

So understood, this model describes the regulation of cyberspace as well. There, too, we can describe four modalities of constraint.

Law regulates behavior in cyberspace--copyright, defamation, and obscenity law all continue to threaten ex post sanctions for violations. How efficiently law regulates behavior in cyberspace is a separate question--in some cases it does so more efficiently, in others not. Better or not, law continues to threaten an expected return. Legislatures enact, prosecutors threaten, courts convict.

Norms regulate behavior in cyberspace as well: talk about democratic politics in the alt.knitting newsgroup, and you open yourself up to "flaming" (an angry, text-based response). "Spoof" another's identity in a "MUD" (a text-based virtual reality), and you may find yourself "toaded" (your character removed). Talk too much on a discussion list, and you are likely to wind up on a common "bozo" filter (blocking messages from you). In each case norms constrain behavior, and, as in real space, the threat of ex post (but decentralized) sanctions enforce these norms.

Markets regulate behavior in cyberspace too. Prices structures often constrain access, and if they do not, then busy signals do. (America Online (AOL) learned this lesson when it shifted from an hourly to a flat-rate pricing plan.) Some sites on the web charge for access, as on-line services like AOL have for some time. Advertisers reward popular sites; on-line services drop unpopular forums. These behaviors are all a function of market constraints and market opportunity, and they all reflect the regulatory role of the market.

And finally the architecture of cyberspace, or its *code*, regulates behavior in cyberspace. The code, or the software and hardware that make cyberspace the way it is, constitutes a set of constraints on how one can behave. The substance of these constraints varies—cyberspace is not one place. But what distinguishes the architectural constraints from other constraints is how they are experienced. As with the constraints of architecture in real space—railroad tracks that divide neighborhoods, bridges that block the access of buses, constitutional courts located miles from the seat of the

government-- they are experienced as conditions on one's access to areas of cyberspace. The conditions, however, are different. In some places, one must enter a password before one gains access; in other places, one can enter whether identified or not. In some places, the transactions that one engages in produce traces, or "mouse droppings," that link the transactions back to the individual; in other places, this link is achieved only if the individual consents. _In some places, one can elect to speak a language that only the recipient can understand (through encryption); in other places, encryption is not an option. Code sets these features; they are features selected by code writers; they constrain some behavior (for example, electronic eavesdropping) by making other behavior possible (encryption). They embed certain values, or they make the realization of certain values impossible. In this sense, these features of cyberspace also regulate, just as architecture in real space regulates.

These four constraints--both in real space and in cyberspace--operate together. For any given policy, their interaction may be cooperative, or competitive. Thus, to understand how a regulation might succeed, we must view these four modalities as acting on the same field, and understand how they interact.

The two problems from the beginning of this section are a simple example of this point:

(a) Zoning Speech. – If there is a problem zoning speech in cyberspace, it is a problem traceable (at least in part) to a difference in the architecture of that place. In real space, age is (relatively) self-authenticating. In cyberspace, it is not. The basic architecture of cyberspace permits users' attributes to remain invisible. So norms, or laws, that turn upon a consumer's age are more difficult to enforce in cyberspace. Law and norms are disabled by this different architecture.

(b) Protecting Privacy. – A similar story can be told about the "problem" of privacy in cyberspace. Real-space architecture makes surveillance generally self-authenticating. Ordinarily, we can notice if we are being followed, or if data from an identity card is being collected. Knowing this enables us to decline giving information if we do not want that information known. Thus, real space interferes with non-consensual collection of data. Hiding that one is spying is relatively hard.

The architecture of cyberspace does not similarly flush out the spy. We wander through cyberspace, unaware of the technologies that gather and track our behavior. We cannot function in life if we assume that everywhere we go such information is collected. Collection practices differ, depending on the site and its objectives. To consent to being tracked, we must know that data is being collected. But the architecture disables (relative to real space) our ability to know when we are being monitored, and to take steps to limit that monitoring. * * *

I noted earlier the general perception that cyberspace was unregulable--that its nature made it so and that this nature was fixed. I argued that whether cyberspace can be regulated is not a function of Nature. It depends, instead, upon its architecture, or its code. Its regulability, that is, is a function of its design. There are designs where behavior within the Net is beyond government's reach; and there are designs where behavior within the Net is fully within government's reach. My claim * * * is that

government can take steps to alter the Internet's design. It can take steps, that is, to affect the regulability of the Internet. * * *

[R]egulability * * * depends upon the architecture of the space, and * * * this architecture can be changed. * * * The code of cyberspace might disable government choice, but the code can disable individual choice as well. There is no natural and general alignment between bottom-up regulation and the existing architecture of the Internet. Enabling individual choice may require collective modification of the architecture of cyberspace, just as enabling collective choice may require modification of this architecture. The architecture of cyberspace is neutral; it can enable or disable either kind of choice. The choice about which to enable, however, is not in any sense neutral.

[A]rchitectures of cyberspace can enable or disable the values implicit in law; law, acting on architectures in cyberspace, can enable or disable the values implicit in code. As one displaces the other, a competition could develop. Authors of code might develop code that displaces law; authors of law might respond with law that displaces code.

East Coast Code (written in Washington, published in the U.S. Code) can thus compete with West Coast Code (written in Silicon Valley, or Redmond, published in bits burned in plastic). Likewise authors of East Coast Code can cooperate with authors of West Coast Code. It is not clear which code one should fear more. The conflict displaces values in both spheres, but cooperation threatens values as well. * * *

This conflict between code and law should push us to consider principle. We should think again about the values that should guide, or constrain, this conflict between authorities. * * * [C]yberspace is not inherently unregulable; its regulability is a function of its design. Some designs make behavior more regulable; others make behavior less regulable. Government * * * can influence the design of cyberspace in ways that enhance government's ability to regulate. * * *

* * * Judge Easterbrook argued that there was no reason to teach the "law of cyberspace," any more than there was reason to teach the "law of the horse," because neither, he suggested, would "illuminate the entire law." This essay has been a respectful disagreement. The threats to values implicit in the law--threats raised by changes in the architecture of code--are just particular examples of a more general point: that more than law alone enables legal values, and law alone cannot guarantee them. If our objective is a world constituted by these values, then it is as much these other regulators--code, but also norms and the market--that must be addressed. Cyberspace makes plain not just how this interaction takes place, but also the urgency of understanding how to affect it.