

What Things Regulate Speech

Lawrence Lessig*

In 1995, California passed a statute making it a crime to sell porn in vending machines. More precisely, the statute made it a crime to sell “harmful matter” (meaning harmful to minors) in any vending machine, unless that vending machine is equipped with an adult identification number system.¹ What an adult identification number system in a vending machine would be, no one quite knows. None exists at the moment, though presumably, a credit card would suffice. Or tokens sold by news agents — at least if sold by a vendor who checked the age of the purchaser.

The aim of the statute is obvious. It is to keep kids from porn. An unattended vending machine can’t tell whether its vendee is 18 or 80. So an unattended vending machine can’t discriminate in its distribution of porn. Porn shouldn’t be distributed by nondiscriminating technologies — or so the California legislature thought. And vending machines are such a technology.

Free speech activists challenged this statute under the First and Fifth Amendments. Their claims were familiar. The statute, they said, reached too broadly. Its effect reached beyond its effect on kids. The law effectively banned porn distributed in this medium (since adult identification systems would be too expensive), and such a ban would limit the rights of adults as well as kids.

There was of course something to this claim. The aim of the statute may have been kids, but its effect was certainly much broader. Its effect would be to require than only humans sold porn. But forcing humans to sell porn creates two different sorts of constraints that constrain adults as well as kids.

One is the constraint of social norms. Norms frown, or better, sneer on porn consumers. Many consumers — call them wimps, or the well-adjusted, you pick — would therefore prefer to purchase porn anonymously. And some of these would only purchase porn if they purchased it anonymously. A statute such as California’s dis-

*Professor of Law, Harvard Law School. Thanks to Melanie Glickson for exceptional research support. This is very much a first draft. Comments would be grateful. Please do not quote without permission.

¹ California Penal Code § 313.1(c)(2), and (h).

criminate against such people. It effectively “abridges” their ability to read this constitutionally protected speech.

The other is the constraint of poverty. Porn (in real space at least) costs money. If porn could be distributed in machines, the cost of distribution may well be lower. Perhaps not much lower, but for the poor, the marginal difference could be quite significant. By eliminating this form of distribution, California is effectively eliminating a particular kind of porn — namely, poor-persons’-porn. And so again, with respect to these people, the law effectively “abridges” the access of these adults to constitutionally protected speech.

But these constraints notwithstanding, the 9th Circuit upheld the statute.² The interest in protecting kids won out; the power of the state to ban the sale of porn in unattended vending machines was upheld. The plaintiffs appealed the case to the Supreme Court. On March 17, 1997, the Supreme Court denied cert in the case, thereby leaving the California law standing.

There was a special irony in the Court’s denial of cert that very week. For the week of March 17th was an important week for technologies that distribute porn anonymously. On Wednesday of that week, the Court heard arguments on the CDA — Congress’s own attempt (failed and stupid that it was) to limit the anonymous distribution of porn. There were of course big differences between the two laws. One difference is a difference in the technology regulated. A second is the language used to pick out the speech to be regulated. But both laws dealt with technologies that made porn accessible to kids. Both dealt with technologies that (in their present state) can’t easily discriminate in the distribution of porn to kids. And both created incentives to modify the technologies to enable them to discriminate on the basis of age. Yet in *Reno v. ACLU*,³ the Court struck down Congress’ statute, while that week it upheld California’s statute without even holding it to consider whether the decision in *Reno* might raise doubts about the decision in *Crawford*.

I set these two cases next to each other not because I think the issues in each are identical. Much separates the two statutes, and little can be inferred from the denial of cert. But the contrast is a reminder of the burdens that first amendment law allows when the Court perceives the burdens to be the only means available to protect kids. *Crawford* might be an extreme, but I suggest it is closer to the

² *Crawford v. Lungren*, 96 F3d 380 (1996).

³ 117 S.Ct. 2329 (1997).

norm than *Reno*. It stands for the rule that has governed in this area since time immemorial — that when kids are on the table, the question is not whether the requirement is too burdensome on free speech, but whether it is greater than it needs to be. The question is whether there is some less burdensome way to achieve the same censoring end. If there is not, the law will stand.⁴

The success in *Reno* came in convincing the Court that other less restrictive means for keeping kids from porn were still available — alternatives to CDA that were less burdensome on speech. More importantly, the success was to convince the court to err on the side of activism. There was little in the Court's past that suggested it would tend to such attention. Historically the Court has been slow to get the significance of a new technology. Historically it has allowed extensive regulation early on, only later cutting back on regulatory power.⁵ Thus it is a testament both to the power of the net, and to the amazing work of groups like EFF, CDT, the ACLU: That within a period of two years, our whole culture could be so infused with this amazing reverence for the net that it could overcome the traditional reluctance of the Court to give full First Amendment protection to a new technology. When the history of the first amendment in this last third of the century is written, these will be real heroes in the story.

But this success just raises new responsibility. The movement has an authority, and it needs to exercise that authority responsibly. The movement must think through the consequences of its fight. It must think through the consequences of the regulatory strategies it is seen to be supporting. These first moves in this regulatory struggle are critical, and they will set a direction that later on can't so easily be controlled.

My sense is that this first major victory — in *Reno v. ACLU* — has set us in a direction that we will later regret. It has pushed the “problem” of kids and porn towards a “solution” that will (from the perspective of the interest in free speech) be much worse. The “less restrictive means” touted by free speech activists in *Reno* are, in my view, far more restrictive of free speech interests than a properly

⁴ This point is made well in Eugene Volokh, *Freedom of Speech and Shielding Children* (unpublished, 1997).

⁵ This was the history of regulation of movies, and television, and radio as well. See Thomas G. Krattenmaker & L. A. Powe, Jr., *Converging First Amendment Principles for Converging Communications Media*, 104 **Yale L.J.** 1719 (1995).

crafted CDA. And unless we quickly shift ground, we will see Congress embracing these less protective means.

My aim in this essay is to demonstrate just why this is so. My argument in the end is that the only constitutional strategy that Congress can follow for regulating “indecentcy” on the net is a strategy very much like the CDA.

I begin, however, a couple steps back. In the section that follows, I offer a general way to think about “vending technologies,” and a specific way to link those technologies to the question of free speech. Against this background, I then sketch the strategy implicit in CDA-like solutions to the problem of indecentcy, and then the strategies offered in CDA’s stead. My claim will be that these alternatives are far more threatening to free speech interests, and (more daringly) that it would be unconstitutional, under present free speech law, for Congress to use its power to advance these alternatives. If Congress is to regulate, its only power is to regulate in a CDA-like way.

In the present climate, of course, this is an unsafe position to take. Unsafe, because the fury of the cyber revolution is quite well advanced. The struggle over defining what cyberspace will be has the feel of the French Revolution. People are shocked at the tone of the debate, terrified at the fury. And in such a context, one is well advised not to step out of line.

Promoting a CDA-like solution to the “problem” of indecentcy is very much to step out of line. And so let me be clear about a couple points up front (not that I think it will matter to Robespierre). I am not advocating a CDA-like solution because I believe there is any real problem. In my view, it would be best just to let things alone. But if Congress is not likely to let things alone, or if the President is more likely to bully a “private solution” than leave things alone, then we need to think through the consequences of these different solutions. We need to compare the consequences, from the perspective of free speech, of adopting one regime over the other. We may well prefer that nothing be done. But if something is to be done, then whether through public or private regulation, we need to think through its consequences for free speech.

THE TECHNOLOGIES OF VENDING MACHINES, VIDEO AND OTHER.

Machines vend. Think about that for a second. If there were a single fact about modern society that would seem most bizarre to citizens of two centuries ago, it would be this. Structures — machines — exist for facilitating market transactions, automatically.

The examples are endless — Coke machines, pay-TV, long-distance telephones, machines selling condoms, television, gas station pumps, name just a few.

The ways of this vending differ, but for our purposes, we can distinguish two patterns along two different dimensions. One distinction is between *push* and *pull* vending; the other is between *discriminatory* and *nondiscriminatory* vending. Let me explain each.

Push vending is vending to the couch potato. Its strategy is to spew forth a string of vending opportunities, and hope that some stick. Television advertising is the best example of this. People watch TV; products are paraded before them; and the hope is that this parade will affect their consumption behavior at sometime in the future. The picture is of the consumer who wouldn't know better; who, but for this spur, would buy nothing, or maybe buy something else.

Pull vending is just the opposite. We imagine the consumer knows what he or she wants. We just make that available, and the consumer will buy it. The coke machine is a simple example. The machine stands there politely, waiting to serve. Someone is thirsty, and comes up to buy what the machine has to offer.⁶ In this sense, the technology just makes things available; the consumer must come and buy what is available.

Discriminatory vending is vending that is in some way conditional — vending only if some condition is met. If you deposit \$1.00, you can have a Pepsi. If you don't, you can't. *Nondiscriminatory* vending is unconditional, vending whether or not a condition is met. By discriminatory, I obviously don't mean anything pejorative. If I only sell books to people who give me money, I am, in the sense I mean, discriminating. On the whole, there is nothing wrong with that discrimination. (I'm told it's the sort of discrimination that makes the world go round.) But nonetheless, it is discrimination, and it requires a kind of technology for achieving its effect.

The breakthrough for vending was such technology — technologies that is that could discriminate. From the beginning of time,

⁶ Obviously, the line between push and pull is not sharp. For example, think about the home-shopping network — the couch potato watches a string of second rate products, and then when one comes that he or she wants, he or she calls a special number and buys it. Is this push or pull? Or think of modern coke machines — huge, and well-lit, more like billboards than boxes, pleading and flashing the imagine of what you should want, wanting you then to come buy from that machine. Is this push or pull?

people had the ability to make non-discriminatory vending machines — boxes, for example, that say “take one” but which could neither collect money for the one taken, or make sure that only one was taken. But when machines could do the discriminating — when the discrimination didn’t need a person — then the opportunities for vending increased dramatically. Vending now was not tied to people. There could be vending in remote places, as well as vending at remote times.

There are now many kinds of discriminations that technology makes possible. Money is just one kind. Time is another: When MCI charges less because you’re calling after 7pm, that too is a kind of discrimination. It, like the ability to collect coins, makes it easier for MCI to vend telephone calls, since it makes it easier for MCI to discriminate in pricing. Identity is a third kind of discrimination: When MCI discriminates on the basis of whether you’re calling a “friend or family” member, the system makes possible a discrimination in pricing that makes for optimal pricing.

Identity discrimination is the future in vending technology. It is also the technology with most to learn. ATMs use passwords as ways to know that the customer is the owner of the account, but passwords can be lost. The NYU Law School library uses a hand-reading machine to vend access to the library, but the technology sometimes acts up. Security companies are beginning to implement technologies that can read a face, or scan an iris, as a way to verify the identity of a person who wants access to a particular place. And keyboard manufacturers have begun to discuss thumb scanners, to attach to keyboards, to verify the person who is using the keyboard. All of these technologies are no doubt imperfect. But then again, no technology for discriminating is perfect. There are slugs in every Coke machine, and there are mistakes in every system of price discrimination. The question is not perfection, but whether the gain from a technology is greater than the cost, where costs includes costs due to error.

We can map these four possibilities with the following matrix, and this matrix will help us categorize the technologies of vending

	<i>Push</i>	<i>Pull</i>
<i>Non-Discriminating</i>	[1]	[2]
<i>Discriminating</i>	[4]	[3]

that exist just now:

Televisions (circa 1965) are box 1 technologies — they are push technologies, which are non-discriminating in the access they grant.⁷ Anyone watching television is exposed to the opportunities TV vends. Newspaper boxes offering free newspapers are in box 2 — pull technologies which are also non-discriminating in the access they grant. Anyone (the rich as well as the poor, the old as well as the young) can open the box and get the newspaper. Coke machines are box 3 technologies. Only those with 75 cents get access to the coke, and these people get to select (pull) precisely what they want. Televisions with cable are box [4] technologies — the discrimination here is the discrimination between those rich enough to afford cable, but what is fed across the cable is not itself selected.

As these examples clearly suggest, a particular technology (a television, or a radio, or a newspaper box) can shift from one box in this matrix to another. One can make a coke machine (in recreation room in one home) non-discriminatory; or one can make public bathrooms discriminatory. Nothing fixes a technology to one of these four boxes. Alternatives exist. And what determines which alternative makes sense is a combination of the costs of different technologies, and the values that different vendings are to embrace.

For most goods, it is enough simply to let the market. Most would have thought it necessary to discriminate in the sale of software; Netscape showed that wasn't quite the right. Most newspapers are for sale, but some are given away. The decision whether discrimination in vending makes sense is a complex judgment about the costs and benefits of discrimination. And when only private values are involved, it is obviously best to leave such decisions to private judgments.

But sometimes public values overlay these private judgments. Sometimes they compete. It might be cheapest to vend condoms at drug stores, but if one constraint on condom sales is the embarrassment of the purchaser, then more condoms might be sold if they were sold in vending machines in bathrooms. The cost of these machines, however, may be too expensive. So in some cases it may

⁷ In a trivial sense, of course, they are discriminatory — you must turn the TV on, so it is in that sense conditional on something. More significantly, it is conditional upon your having a television set. These qualifications are all correct, but unnecessary for the purposes of this essay. The boxes here have fuzzy borders, and it is not essential to find paradigm cases.

make sense for a public that wants more condoms used to subsidize machine vending.

Or again: It might be cheapest to vend cigarettes in machines. But since cigarette machines can't discriminate on the basis of age, a public policy against the sale of cigarettes to minors might direct that vending machines not be used.

Or finally: The same with whiskey. Whiskey might be optimally sold in vending machines. But uncontrolled access to whiskey may be inconsistent with other public values. These values then may direct that machines not be used to sell this alcohol. Sometimes public values alter the optimal vending strategy, or better, the optimal vending strategy given public values is often different from the optimal vending strategy without public values.

With commodities, there is little to limit the public's imposition of an alternative vending strategy. For most products, any burden is OK. There is no constitutional problem with a law making it illegal to sell cigarettes by vending machines. Nor a law that makes the sale of liquor illegal except in state stores. The constitution has been read not to give merchants of stuff any special protection against the discriminatory regulations of the state. The state is allowed to discriminate in all sorts of ways to make sure that products are sold only in certain places, and only to certain people.

But when it comes to speech, the rules change. The state is not as free to determine the type of vending that shall be allowed. For most speech, the constitution requires that the state permit whatever type of vending individuals want. The power to regulate that vending is constrained by the first amendment. And for most speech, this constraint is real.

Rules limiting the Rules for Vending Speech

For our purposes, we can understand free speech law to divide speech into three classes. One class is speech that everyone has the right to. For this class, the state's power is quite slight. The paradigm here is political speech, but in effect it includes any speech not described in the next two classes.

A second class is speech that no one has the right to. Here the state's power is practically unlimited. The paradigm here is obscene speech, or more strongly, child pornography.

And a third class is speech that people over the age of 17 or 18 have a right to, while people under those ages do not. This we can call "indecent" speech, which the state has regulated in many differ-

ent contexts. *Ginsberg v. New York*⁸ is the paradigm case. New York had a statute that banned the sale of certain speech⁹ to anyone under the age of 17. Implied in this rule was the requirement then that vendors in New York check the ID of people who wanted to purchase such porn. But because this burden seemed relatively slight, and no cheaper discrimination seemed possible, the Court found the burden constitutionally permissible.

The essence of the state's power over this third kind of speech is the power to zone speech on the basis of age. Zoning here simply means conditioning access on the basis of some criteria—a state imposed discrimination, for selected features of the person or speech. With “indecent” speech, then, the state has the power to mandate discrimination in the distribution of such speech. In general, it has the power to require discrimination, so long as the discrimination required is the least restrictive form of discrimination that existing technology permits.

Crawford is the extreme in this line. Like the statute in *Ginsberg*, the aim of the California statute in *Crawford* was to require discrimination in the sale of porn. But unlike *Ginsberg*, to implement this requirement would cost real money. Indeed the requirement in *Crawford* is so great that it essentially means that it is impossible to continue to sell porn from vending machines. But that fact was not enough to strike the statute. There was no better way to protect kids, and so even if the consequence was a burden imposed on adults, the statute would withstand constitutional review.

One should not of course place too much weight on *Crawford*. It is an opinion of the 9th circuit, and the court's denial of cert is not an endorsement of its result. But the opinion is not out of step with the general run of cases applying this indecency principle. In each case, the question is not whether the burden is in some absolute sense too high. The question is rather whether the burden is higher than it has to be.

⁸ 390 U.S. 629 (1968).

⁹The statute in *Ginsberg* defined “harmful to minors” to mean “that quality of any description or representation, in whatever form, of nudity, sexual conduct, sexual excitement, or sado-masochistic abuse, when it: (1) predominantly appeals to the prurient, shameful, or morbid interest of minors, and (2) is patently offensive to prevailing standards in the adult community as a whole with respect to what is suitable material for minors, and (3) is utterly without redeeming social importance for minors.” *Id.* at 644. The Court characterized this test as the obscenity test as applied to minors. See *id.* at 636.

The same explains the Court's television indecency cases. When television and radio were nondiscriminating push technologies, they did not allow ex ante control of what came across the channel. One tuned in a television or radio station, and across the channel came what someone else has selected to send. At the time, this was the nature of television. There wasn't a feasible technology to make television discriminatory. Hence there was an unavoidable conflict between the interest of television and radio to broadcast what they want, and the interest in the state, and parents, to block access by kids to certain kinds of speech. The compromise was time-shifting — requiring that content of a certain type be moved to a time when kids are unlikely to be watching television. Television during ordinary hours was made into Disney World, while television during evening hours was relatively free.

This was no small burden on adults. But the compelling interest in protecting kids, and the relatively fixed technology, meant that this shifting was the only solution that would satisfy the state's compelling interest.

The V-chip changes this. The V-chip permits a kind of automatic blocking of television access based on ratings of television shows. If enabled, the chip would, for example, block a television's viewing of R-rated material. It could block that viewing at the level of the show itself. Or more ambitiously, one might imagine a system where the show itself was modified depending upon the selection in the V-chip. If for example, an R rating were selected, then the system might show more of the sex scenes in a particular show; if PG were selected, it might edit those scenes out.

The essence of this system is to enable viewers with ex ante control over the content of speech viewed. As Jack Balkin describes it, its essence is to change the technology of broadcasting.¹⁰ It is to shift television from a non-discriminating push technology to a discriminating push technology. Turn on the filter, and it will not push material blocked out. Turn off the filter, and it will push all the material there is.

I have questions that we will consider below about the scope of the filter that the V-chip would permit. But assuming it were limited just to *Ginsberg* like speech, it should be plain that the emergence of the V-chip as a possible technology of discrimination undermines the position of the FCC in *Pacifica*. The V-chip becomes a less burdensome technology for discriminating — less burdensome, that is,

¹⁰ J. M. Balkin, *Media Filters, the V-Chip, and the Foundations of Broadcast Regulation*, 1996 **Duke L.J.** 1131 (1996).

on adults. Rather than make all TV fit for the Disney Channel, we can make all TVs capable of translating the broadcast into Disney material.

This is again consistent with the general rule applicable to indecent speech. As Eugene Volokh puts it, the question is not whether the government will be allowed to regulate to protect kids. Always it will be — so long, that is, as it selects the least burdensome technique of regulation. But as an absolute, there is nothing that guarantees that this “least burdensome” means is not itself burdensome. It is always a question asked relative to other techniques. The assurance is just that there won’t be a burden greater than is necessary. That is the essence of this third kind of regulation.

Applying the Rules Limiting the Rules for Vending Speech to Video Vending Machines (a.k.a., computers)

Computers too are vending machines. They are a kind of video vending machine. But the range of vending machines we’ve so far discussed, they are by far the most flexible, or the most plastic. A survey of the net collects examples that fit within each box of the matrix. Consider some examples here.

- Box 1 Some parts of the net web effect non-discriminatory push technology. This is the newest part of the web, and includes technologies such as Pointcast. In these spaces, the web feeds information without any formal or mandatory discrimination. (The user, of course, can select to block certain kinds of information by selecting other information.)
- Box 2 Most of the web is non-discriminatory pull technology. Users search on what the web has available, and then go to those places, and retrieve what they want. Where they go is not blocked by who they are, and what they get is determined by what they want.
- Box 3 Another new area of the web is discriminatory pull technology. The best example of these are zines that charge for access — the Wall Street Journal, for example. But charging is not the only discriminatory technology. Some journals require that users register. That registration then is used to profile use, so that the site can sell advertising. The key here is that access is blocked, and that user select what they want.

Box 4 Finally, there are spaces on the web that discriminate in access, but have push content. Adult sites are the best example here. Users establish an account (usually one they must pay for), and then get access to spaces where content is pushed to them — chat rooms, or video spaces, where the user, like with television, can sit passively at the machine, as the machine feeds content.

As a class, machines on the net span all four boxes of my matrix. Or put differently, the architecture of the net can be designed to fit a vending technology within any of these four boxes.

More importantly, however, is the difference that there is in the cost of changing these technologies from one box to another. Relative to the cost of real space code, the cost of altering the code of cyberspace to alter the technology of vending is extremely low. To move a product from one box to another requires simply a change in code — bits, organized in software.

This plasticity is both an opportunity and a threat. It is an opportunity because it means that the state's legitimate interest in regulating certain kinds of speech, by mandating certain kinds of discrimination, can be effected at a relatively low cost. And if it is true that these interests in the past have always justified state regulation, then for these topics of regulation, we might expect the burden of the regulation to fall. The same state interests will be advanced, but at a cheaper cost.

But this plasticity can also become a threat. For if there is a cost with discrimination in the market for speech, then the cheapness of technologies for discriminating threatens to increase the incidence of this discrimination. The threat is that the government will leverage its legitimate interest in pushing some discrimination into an interest in pushing technologies that discriminate more generally.

We can see the point most plainly in the context of proposals for dealing with "indecenty" on the net. All of the proposals would facilitate discrimination in the distribution of porn on the net. All are technologies for moving porn from box 2 to box 3. What distinguishes the proposals are (1) how expensive each is, and (2) how general the discrimination each facilitates is. Put another way, some make possible a general system of discrimination, while others make possible discrimination in only some areas.

My argument is that the government has no legitimate interest, consistent with the First Amendment, in facilitating general discrimination; that the most the First Amendment can permit are regulations that push the net to discriminate in narrowly drawn areas.

This is not to argue that it would be unconstitutional if the net became a place where general discrimination were possible; it may well become that place, but that's a different point. My claim is only that the government's role in facilitating this generalized discrimination is quite limited, and that we should select strategies for advancing its legitimate interests that don't easily generalize to this broader power.

The general question we should ask is how much speech gets vended in box 3 structures, and how much control do these box 3 structures give to governments and institutions of censorship. My claim will be that among all the alternatives that have been suggested for dealing with the "problem" of indecency, a CDA-like solution would minimize the amount of speech subject to discriminating technologies, or at least minimize the role the government has in facilitating this discrimination. Or again, that among the solutions proposed, a CDA like solution is most speech protective.

The Regulation of the CDA

First, however, we must distinguish two issues when considering the CDA — one is the scope of speech covered, and the second is the way in which the regulation had its effect. There is no doubt that because of Congress' stupidity with respect to the first question, CDA was unconstitutional. Its definition of the speech covered was are too vague to pass constitutional review. And it seemed plainly to extend beyond the scope of *Ginsberg* speech.

But my interest in considering the CDA is tied to the second issue — the way in which the regulation was to have its effect. For the statute functioned first by banning a certain kind of speech unless second, that speech was put behind a wall that with "reasonable effectiveness" screened out kids. The technique was not filtering. It was to set up identity checks on the doors through which people wanted to pass.

The scheme depends upon a system of adult identification numbers. The statute gave a defense to a publisher who used such a system to zone his space into spaces where kids couldn't come. The system didn't have to be perfect — it had simply to be "reasonable effective" in keeping kids out. And one fair reading of "reasonable" — indeed the reading that all of torts law give to same word — would be "reasonable given the technologies that exist."¹¹

¹¹ It is therefore a bit odd that the Court in *Reno* stuck firm in its reading of "effective" to conclude that the statute was too demanding. Why, rather than striking the statute, or any similar statute, it couldn't have read "reasonably effective" as a negligence standard is unclear.

When the government argued this case, they either had no idea about how the technology worked, or they had no interest in winning the case. For they stipulated to facts which were not then, and certainly are not now, true.¹² The picture they had, and that the ACLU was happy to agree to, was that each site would have to run its own identification number system. And they conceded that the costs of running such a system would be quite high. Nonetheless, they argued, that burden was well worth the benefit of keeping kids from porn.

Even here, the argument should strike one as absurd. In real space, we have all sorts of places where IDs must be checked — bars, convenience stores, college parties, airports. But no one thinks that the obligation to check IDs means the obligation to create an ID system specific to that place. People don't have an ID for their local pub, and different one for their local grocery store. Instead, in real space, standard IDs have developed — drivers licenses, for example — which stores can rely upon to check age.

There is no reason to think the same can't develop in cyberspace, every reason to believe it would, and all the evidence to suggest that it already has. The net is filled with ID companies which will, for a fee, issue an ID which is then useable at any number of places to check the age of participants. These ID systems are relatively cheap, and given the low cost of net transactions, their cost is likely to fall even more.

But the Court hesitated before embracing this picture of an ID enabled cyberspace. Its hesitation was a function either of the poor state of the record, or a limit in imagination. For not only should we expect that the market in cyberspace will provide IDs, we should expect that will provide IDs more efficiently the market in real space provides IDs. As it is right now, it takes but a few minutes to get a cyber-ID. There are a range of places one could get them. They are

¹² Two examples should suffice. In finding 96, the court found that “content providers who publish on the World Wide Web via one of the large commercial online services, such as America Online or CompuServe, could not use an on-line age verification system that requires cgi script because the server software of those on-line services available to subscribers cannot process cgi scripts.” *ACLU v. Reno*, 929 F.Supp 824, 845 (1996). But why? There is nothing magic about making code that can process cgi scripts. The same with the findings regarding age verification technologies. Finding 90 reports that there is “no effective way to determine the identity or age of a user who is accessing material through ... newsgroups.” Well again, that depends on the software. Imagina Corporation, for example, sells a NNTP server that does allow for identity verification.

not all associated with adult sites.¹³ And finally, we could envision sites that could credibly guaranteed privacy protection with the use of their IDs.

The state, for example. Imagine in a new CDA (CDA II) that Congress simultaneously established an online ID issuing authority, or better, licensed private companies to serve as online ID issuing authorities. Imagine these issuing authorities issued pseudonymous IDs for a small fee (say, \$2 an ID), and were required to protect the privacy of the ID holders by never disclosing the pseudonym's holder. Sites then could link to one of these sites, and verify access of anyone who wants entry.¹⁴

This is the picture of a CDA that I want to consider as the baseline against which the alternative will be measured. Its features are these: First, its restriction extends only as far as the constitutionally legitimate governmental interest. The statute functions as a zoning statute, but it does not enable zoning any more extensive than its narrow interest. Second, its restriction will not generalize into a more comprehensive system for filtering or blocking speech. It is a targeted blocking system, not a generalizable one. Its burden would be that those who engage in *Ginsberg* speech must do it in a context where other have been screened.

Now I don't mean to say that there won't be hard questions even with a technology as minimally intrusive as I have described. It seems to me that under the *Ginsberg* standard there are easy cases on both sides, and hard cases in the middle. The easy cases are porn sites (which would be within the range of legitimately regulated speech) and sex education or health sites (which would be within the range of unregulable speech). In the middle there will be many contexts which the Court will have to resolve. Adult chat rooms? Or general chat rooms where indecent language is used? In my view, none of this should be found to be regulable, but that's a battle to be fought later on. For now, the important idea is simply the structure of this regulation.

How does it compare with the alternatives?

ALT.FILTER

¹³ So concern that one must be seen to associate with adult sites need not be a problem.

¹⁴ There are more elegant ways to achieve the same end, but this is sufficient for now.

The alternatives to the CDA are all what we might call “filtering” solutions. They all facilitate a kind of content filtering (not identity blocking). And so we should begin by thinking a bit more carefully about what it means to live in a *filterable* world.¹⁵

There alternatives are essentially two. The first is an earlier version of the second, but I will discuss it in any case because it is the version that will be litigated first. This is private blocking software — software like CyberSitter, or SurfWatch. The second alternative is a far more general and powerful filtering or content selecting standards developed by the World Wide Web consortium. This is PICS. In the section that follows, I will quite briefly discuss the problems with blocking software. In the section following that, I will turn to PICS.

The Bad in Blocking

Private blocking works like this: Companies compete to gather lists of sites on the web that “parents” don’t want their “kids” to go. The companies advertise the kinds of sites that get on their lists. Some build lists focused just on sexual material. Others build lists focused on violence. Others build lists focused on sites that send the wrong message about sex, or smoking. The lists are lists of banned books, but the actual list of banned books is never published. It, instead, is a trade secret. For if it were published, then it could be copied, and then the value in constructing the list in the first place would be lost.

So instead, the list is encrypted, and delivered on a regular basis to purchasers of the software. The software itself costs around \$50.00; updates can cost between \$10-\$20 a cycle.¹⁶

Now the model for this software sounds good enough — those who need filtering of the web buy it; those who don’t don’t have suffer the burden of kids. And because individuals can select among a range of companies and range of products, it is as if individuals are simply selecting their censor. It is the same sort of thing that I do when I select one newspaper rather than another, or when I subscribe to one cable channel rather than another.

¹⁵ For an early, and excellent, analysis of the same issue, see Jonathan Weinberg, *Rating the Net*, 19 **Comm/Ent** 453 (1997). Weinberg’s analysis emphasizes a balance in the question of the costs and benefits of rating. My analysis emphasizes who the technology empowers.

¹⁶ See, e.g., <http://www.cybersitter.com/cysitter.htm> (Cybersitter, which has free updates); <http://www8.zdnet.com/pcmag/features/utility/filter/ufu1.htm> (Cyber Patrol, which has a \$20 subscription price).

The technology thus shifts the architecture of the net (for those who use the software) from a box 2 technology to a box 3 technology — from a non-discriminating pull technology to a discriminating pull technology. But the tools of the discrimination here are extremely crude, and opaque. One can't know what's on these lists, and there's no simple way to verify that sites are not included for the wrong reasons. Horror stories already abound — sites opened to criticize this blocking software themselves included in the blocked list. Sites opened to discuss AIDS, or gay rights, excluded because of mistaken associations with indecency. Controversial sites are easily excluded, yet there is no way to know just who gets cut.

This is the first and greatest problem with a regime of private blocking. For contrary to the suggestion often made about governmental censorship, we should expect far more censorship when the censorship is private, rather than public. The reason is a simple point of economics. It costs money for the government to prosecute a particular site, and that prosecution is supervised by reviewing courts. It costs nothing for a private company to block a particular site, (especially if the exclusion is opaque) and there is no system of review for that blocking.

From the perspective of free speech interests, this problem is bad enough. But when the government buys into this method of blocking — when libraries, for example, begin to use this software — this tension with free speech interests transforms into a conflict with constitutional values.

Take the case of a public library installing these blocking systems on their internet links. The first question should be whether this blocking software is on all machines, of just machines that children have access to. There would no justification for a public library installing such filtering generally. If there is a justification at all, it must be limited to kids.

But even as to kids, these lists are problematic. For again, they filter far beyond the constitutional justification for governmental filtering. They are not filtering on the basis of *Ginsberg* speech; they are filtering on the basis of what the market wants. Kids are blocked by state action from speech they have a constitutional right to view.

The Worse in Labels

As bad as private blocking is, however, it does have its virtues. I've argued that its aim is to make some speech subject to box 3 filtering. But its virtue is that it leaves the balance of speech in box 2. Its aim is not to make all speech subject to discriminating technologies. It targets that technology to just some kinds of speech.

The second alternative to CDA is not so discriminating. This is the technology of PICS. PICS is a more efficient long term solution to the problem of filtering than blocking software — it is cheaper and more general and more open to competition. But its consequences for free speech are more dramatic as well.

To see why, we need once again to step back from the question. There was a certain rhetoric that went with the founding of cyberspace, a rhetoric which now has become unthinking commonplace. The commonplace goes something like this: Cyberspace is an unregulable space. It is a space where the cost of exit is extremely slight. Because so slight, any burdens imposed by a central authority are burdens that are cheaply routed around. In one of the clichés of the e! generation — the net interprets censorship as failure, and simply routes around it.

In area of content regulation, one reason the net was unregulable was that content was so hard to identify. The best machines in the world couldn't distinguish an obscene short story from a sex education text book, or a skin-zine from a medical text. Theorists told us that automatic filtering was computationally impossible. And this impossibility in turn was the ground of our freedom. It couldn't be done, and therefore we didn't need to filter it.

But one of cyberspaces' most important theorists had, early on, a caveat to add to this story — a caveat in the form of a warning. For while it is true that with the present architecture of the net, machines couldn't censor, Nicholas Negroponte warned that it would take a tiny change to erase this impossibility. For if material on the net were labeled, then the dumbest machines on the net could filter. Once labeled, the problem of filtering becomes quite simple. And hence to enable censorship, Negroponte warned, governments would only have to enable labeling.

PICS is a system, and the incentive, for enabling that labeling. Its idea is simple, and ingenious. The developers realized that there were two conceptually distinct parts to the problem of filtering on the net. One is the problem of software to actually implement any filter. The other is problem of rating the content on the net, to make any filter workable. W3C decided to separate the two problems of filtering and labeling. As a recent FAQ on PICS and Intellectual Freedom describes it,

A label describes the content of something. A filter makes the content inaccessible to some audience. While both labeling and filtering may introduce social concerns, the concerns are somewhat different. More generally, there are six roles that could all be filled by different entities:

1. Set labeling vocabulary and criteria for assigning labels
2. Assign labels
3. Distribute labels
4. Write filtering software
5. Set filtering criteria
6. Install/run filtering software

PICS itself actually fills none of the six roles listed above! PICS is a set of technical specifications that makes it possible for these roles to be played by independent entities.¹⁷

Now in some ways, the solution of PICS seems ideal. For not only does it enable individuals to select the rating system they want, it empowers individuals or groups to set up competing rating systems. The system is therefore *horizontally* neutral — the Christian Right can have a rating system, as can the Atheist Left. Individuals can then select the rating system that each then wants.

But the system is more than horizontally neutral. It is vertically neutral as well. It not only allows any number of filter to be selecting among, it also allows the filter to be imposed — invisibly — at any level in the distributional chain. The filter can be imposed at the level of the individual's computer. But it can also be imposed at the level of the ISP. Or at the level — in principle — of a nation-state. PICS, as its founders describe it, is “neutral” among these different locations for the imposition of the PICS filter; it has no built in bias for or against this centralized control.

The founders of PICS might be neutral about this control, but we should not be. We should not be neutral about a technology that facilitates state censorship as we as individual censorship, just as we should not be neutral about distributing nuclear bombs to the North Koreans as well as New Zealand. At the very least, it is a dangerous idea (from a free speech perspective) to implement a technology that enables cheap centralized filtering. At a minimum, we should ask in some context where the political implications of this can be measured, whether it is a good thing for us to flip the essential character of the net — not just for us, but for the world — just because we have this obsession with indecency.

And flip the character it no doubt will. Because as has almost become trite, the old architecture of the net resisted just this sort of control. But a PICS enabled architecture facilitates it. It facilitates it not just to deal with indecency — but to deal with any specific content, as well as wide range of other topics. PICS enables filtering on

¹⁷ See <http://www.si.umich.edu/~presnick/pics/intfree/FAQ.htm>

the basis of indecency, or Nazi speech, or criticism of the Chinese government, or questioning of the Singapore parliament. The architecture is scaleable in a way that a CDA architecture is not.

My sense is we should have a debate about this. We should have a debate where neutrality on the question of censorship is not given to us as a technological necessity. We should have a debate about what the character of the net should be, and how we want it to develop. We need to have this debate, long before the defaults get switched, and the system gets turned on.

But my point here today is not so much to argue the virtues or vices of PICS. Its present neutrality on the question of censorship seems to me a vice, but I conceded, that in many other ways, the system has virtues.¹⁸ Whether over all it is system that makes sense is a hard question, and one we should not try to answer in the context of this very specific debate about indecency.

My argument is again not about that question. My argument about the scope of Congress' power. I don't think Congress *can* constitutionally answer that more general debate in the context of a debate about indecency. In my view, that is, it would unconstitutional for Congress to embrace a system like PICS in order to deal with the problem of indecency.

My aim in the last section of this essay is to sketch that argument. In my view the most Congress can do, if in fact it has to do anything, is a revised version of the CDA. If it tries to do more, then it will cross a constitutional line that has been well developed in the past 25 years. The market may develop PICS, and adopt PICS. But Congress can't — consistent, at least with First Amendment values.

The Constitutional Problem with General Filtering

The constitutional problem with PICS can be stated quite simply: PICS would push the architecture of the net from box 2 to box 3. It would push the architecture of the net from a default of nondiscrimination, to a default of discrimination. And it would push this default not just for a narrow class of speech, but for speech quite generally. It would push the net to discrimination across the full range of speech, and it would facilitate this discrimination at any level in the net's distributional chain.

¹⁸ The most obvious virtue is the protections to privacy the system facilitates. Since the system is simply a way to verify assertions, one could verify that a site was privacy protecting by filtering it according to some privacy protection list. The browser would then block me from accessing sites where my privacy wasn't protected.

In my view, this change is far beyond any legitimate interest that the government has in facilitating indecency discrimination on the net. The government's legitimate interest in labels — or in filters, or in blocking access to speech — is quite narrow. And that narrowness should limit the kinds of labeling regimes that the government can, legitimately, support. If the only legitimate reason to block is to block access for speech of kind X, then my it is not constitutional for Congress to require technologies that enable blocking for speech of kinds X, Y, and Z.

There is no clear precedent for this conclusion; just hints from two lines of cases.¹⁹ These two lines of authority suggest the diffi-

¹⁹ One line considers the constitutional of congress' attempt to change a vending technology from push to pull. This is the case of *Bolger v. Youngs Drug Products Corp*, 463 U.S. 60 (1983). At issue was a statute that banned the unsolicited mailing of information about contraceptives. This information was not obscene, and not even *Ginsberg*-obscene (obscene for kids). Nonetheless, the claim was that most would find it offensive; and that a requirement that it only be sent if called for (pulled rather than pushed) was a small burden to avoid this offense. The Court struck this statute. The government couldn't push the distribution of this material into a single form of vending, since the government had no interest in regulating this class of speech. Offensiveness was not a sufficient condition for giving the government regulatory power. Something more was needed. If the market wanted to vend via push, the government could mandate that it vend only via pull. The essence was that the government couldn't interfere to tilt the balance one way or the other. The same point was made much earlier in the case of *Lamont v. Postmaster General*, 381 U.S. 301 (1965). There the question was whether the government could hold at the post office "foreign communist propaganda," and require that the intended recipient return a postcard, requesting that it be delivered. This again was a regulation that was changing the vending of communist propaganda from push to pull. The Court again struck the statute down. The government had no role to play in determining whether the technology would be vended in one way or the other. The first amendment required that it not interfere in the manner of its distribution. The issue with these three technologies for filtering is similar to the issue in *Lamont* and *Bolger*, though importantly distinct. It is similar in that the government would be moving the market from nondiscriminating distribution to technologies that discriminate. It would be tilting the market to facilitating discrimination, and away from nondiscriminating vending. Thus, while this is not a push from push to pull vending, it is a push from non-discriminating to discriminating vending. And one might well conclude that the same first amendment interests are at stake.

The second line of cases is a single point, which admittedly (as with any point) could point in any number of directions. This is the case of *Rowan v. US Post Office*, 397 U.S. 728 (1970). The post office regulations allowed an individual to have the post office order a sender no longer to send material to that individual. It allowed, that is, filtering on the basis of who the sender was. The court upheld these regulations, but on the condition that the government was making

cult that any regulation pushing discrimination is likely to face. But the real test is not old precedent; the real test is what values do non-discrimination advance. Why is a regime that pushed discrimination worse than a regime that did not? What is the cost of discrimination?

This isn't the place (as there isn't the space) to present fully the argument in favor of nondiscrimination. Its essence is that there is a value in the inefficiency of a speech market; that the mix that inefficiency permits is a valuable mix. But for purposes of this debate, the argument is quite straight forward: A regime that facilitated generalized filtering is a regime that forces speakers to self label. You label, or you fall off the screen. By changing the defaults of the system, the regime would increase the costs on all speakers. But for some speakers — indeed for most speakers — the government has no legitimate interest in increasing the cost of their speech. The government can impose the costs of discrimination only on speech that can be legitimately zoned. And most speech cannot legitimately be zoned.

CONCLUSION

Law regulates speech, but not only law. Norms regulate speech, and so too does the market. But the regulator I have tried to focus in this essay is the regulation of architecture — the regulation that gets effected by the very design of a free speech place.

Cyberspace is a free speech place, and its architecture is a regulator. In its early life, it was a space where its architecture — its code — facilitated very little centralized control. It was a space that worked against such control. It was an architecture of relative anarchy; anathema to governmental control.

My argument has been that we consider the alternative architectures that these solutions to the “problem” of indecency provide, and the values implicit in each. We should worry over how effectively an architecture discriminates, and over who gets to effect that discrimination. Among the alternatives I have surveyed here, the solution most likely to succeed (PICS) is also the solution that effects the

no judgment about what kind of speech was being filtered. The government couldn't be charged with making the judgment of “similar” senders. The government could only carry out the recipient's orders. In one sense then one might argue that this is the ratification of government supported discrimination. Indeed, the government is carrying out the order of discrimination — it is ordering the sender no longer to send. But in another sense, this kind of discrimination is extremely ineffective. For it doesn't extend to other like cases. It is limited to only the names the recipient states. The recipient can't say “I want to political mail,” or “I want no mail from Republicans.” If the regulation had permitted that kind of filtering, the Court strongly indicated it would have to strike it down.

broadest discrimination. It is an architecture anathema to the original net. It is a code of control.

We have won the first battle in this free speech struggle. We must now make certain that we don't lose the war. The victory in *Reno* will push Congress to be more careful before it acts again. It might push it not to act again at all. But in this lull, the threats that it will act, and the cajoling of the President to get private interests to act, are now changing the architecture of the net. The threat now is not so much a regulation by Congress; the threat now a regulation by the code. Our attention must be on how the architecture of the net is regulation — what its values are, and what the government's role is in making the values as they will be.

Our tradition is to fear government's regulation, and embrace private regulation. Our intuitions are trained against laws, not against code. But my argument in this essay has been that we understand the values implicit in the code, and that we push for the architecture that best advances free speech values.

My fear is that our tradition, and intuitions, will blind us to the regulatory power that this shift in architectures is bringing about. That unless we understand how code is now a sovereign, we will not see how code can abridge speech.