

CODE

version 2.0

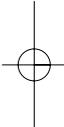
LAWRENCE LESSIG



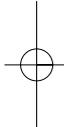
A Member of the Perseus Books Group  
New York



Copyright © 2006 by Lawrence Lessig CC Attribution-ShareAlike



Published by Basic Books  
A Member of the Perseus Books Group



Printed in the United States of America. For information, address Basic Books, 387  
Park Avenue South, New York, NY 10016-8810.

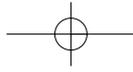
Books published by Basic Books are available at special discounts for bulk purchases  
in the United States by corporations, institutions, and other organizations. For more  
information, please contact the Special Markets Department at the Perseus Books  
Group, 11 Cambridge Center, Cambridge MA 02142, or call (617) 252-5298, (800)  
255-1514 or e-mail [special.markets@perseusbooks.com](mailto:special.markets@perseusbooks.com).

CIP catalog record for this book is available from the Library of Congress.

ISBN-10: 0-465-03914-6  
ISBN-13: 978-0-465-03914-2

06 07 08 09 / 10 9 8 7 6 5 4 3 2 1





# S I X

---

---

## c y b e r s p a c e s

I'VE SAID WE CAN DISTINGUISH THE INTERNET FROM CYBERSPACE. TO MAKE THE distinctive form of regulation that is the subject of this part salient, we need to say a bit more about this distinction. The Internet is a medium of communication. People do things “on” the Internet. Most of those things are trivial, even if important. People pay bills on the Internet, they make reservations at restaurants. They get their news from the Internet. They send news to family members using e-mail or IM chat. These uses are important in the sense that they affect the economy and make life easier and harder for those using the Internet. But they're not important in the sense that they change how people live. It's very cool that you can buy books with one click at Amazon. I buy tons (maybe literally) of books I wouldn't otherwise have bought. But my life has not been changed by one-click (even if my bank account has). It's been made easier and more literate, but not anything fundamentally different.

Cyberspace, by contrast, is not just about making life easier. It is about making life different, or perhaps better. It is about making a different (or second) life. It evokes, or calls to life, ways of interacting that were not possible before. I don't mean that the interaction is new—we've always had communities; these communities have always produced something close to what I will describe cyberspace to have produced. But these cyberspace communities create a difference in degree that has matured into a difference in kind. There is something unique about the interactions in these spaces, and something especially unique about how they are regulated.

Life in cyberspace is regulated primarily through the code of cyberspace. Not regulated in the sense of Part I—my point is not that the code makes it easy to know who did what so that penalties can be visited upon those who behaved badly. Regulated in the sense that bars on a prison regulate the



movement of a prisoner, or regulated in the sense that stairs regulate the access of the disabled. Code is a regulator in cyberspace because it defines the terms upon which cyberspace is offered. And those who set those terms increasingly recognize the code as a means to achieving the behaviors that benefit them best.

And so too with the Internet. Code on the Internet is also a regulator, and people live life on the Internet subject to that regulation. But my strategy in this chapter is to begin with the more obscure as a way to build recognition about the familiar. Once you see the technique applied to worlds you are unlikely to inhabit, you will recognize the technique applied to the world you inhabit all the time.

Cyberspace is not one place. It is many places. And the character of these many places differ in ways that are fundamental. These differences come in part from differences in the people who populate these places, but demographics alone don't explain the variance. Something more is going on.

Here is a test. Read the following passage, and ask yourself whether the description rings true for you:

I believe virtual communities promise to restore to Americans at the end of the twentieth century what many of us feel was lost in the decades at the beginning of the century—a stable sense of community, of place. Ask those who've been members of such a virtual community, and they'll tell you that what happens there is more than an exchange of electronic impulses in the wires. It's not just virtual barn raising. . . . It's also the comfort from others that a man like Phil Catalfo of the WELL can experience when he's up late at night caring for a child suffering from leukemia, and he logs on to the WELL and pours out his anguish and fears. People really do care for each other and fall in love over the Net, just as they do in geographic communities. And that "virtual" connectedness is a real sign of hope in a nation that's increasingly anxious about the fragmentation of public life and the polarization of interest groups and the alienation of urban existence.<sup>1</sup>

There are two sorts of reactions to talk like this. To those who have been in "cyberspace" for some time, such talk is extremely familiar. These people have been on different kinds of "nets" from the start. They moved to the Internet from more isolated communities—from a local BBS (bulletin board service), or, as Mike Godwin (the author of the passage) puts it, from a "tony" address like The WELL. For them the Net is a space for conversation, connections, and exchange—a wildly promising location for making life in real space different.

But if you are a recent immigrant to this “space” (the old-timers call you “newbies”), or if all you do on the Internet is check your stocks or look up movie times, you are likely to be impatient with talk like this. When people talk about “community,” about special ways to connect, or about the amazing power of this space to alter lives, you are likely to ask, “What is this idea of cyberspace as a place?” For newbies, those who have simply e-mailed or surfed the Web, the “community” of the Net is an odd sort of mysticism. How can anyone think of these pages full of advertisements and spinning icons as a community, or even as a space? To the sober newbie, this just sounds like hype high on java.<sup>2</sup>

Newbies are the silent majority of today’s Net.<sup>3</sup> However much one romanticizes the old days when the Net was a place for conversation and exchange, this is not its function for most of its users now. There are exploding communities of bloggers and creativity. But bloggers are still just 3 percent of Internet users; the vast majority of Internet use has no connection to any ideal of community.

Cyberspace has changed in its feel.<sup>4</sup> How it looks, what you can do there, how you are connected there—all this has changed. *Why* it has changed is a complicated question—a complete answer to which I can’t provide. Cyberspace has changed in part because the people—who they are, what their interests are—have changed, and in part because the capabilities provided by the space have changed.

But part of the change has to do with the space itself. Communities, exchange, and conversation all flourish in a certain type of space; they are extinguished in a different type of space.<sup>5</sup> My hope is to illuminate the differences between these two environments.

The next sections describe different cyber-places. The aim is to build intuitions about how to think through the differences that we observe. These intuitions, in turn, will help us see something about where cyberspace is moving.

#### THE VALUES OF A SPACE

Spaces have values.<sup>6</sup> They manifest these values through the practices or lives that they enable or disable. As Mark Stefik puts it:

[B]arriers within cyberspace—separate chat rooms, intranet gateways, digital envelopes, and other systems to limit access—resemble the effects of national borders, physical boundaries, and distance. Programming determines which people can access which digital objects and which digital objects can interact

with other digital objects. How such programming regulates human interactions—and thus modulates change—depends on the choices made.<sup>7</sup>

Choices mean that differently constituted spaces enable and disable differently. This is the first idea to make plain. Here is an example.

At the start of the Internet, communication was through text. Media such as USENET newsgroups, Internet Relay Chat, and e-mail all confined exchange to text—to words on a screen, typed by a person (or so one thought).

The reason for this limitation is fairly obvious: The bandwidth of early Net life was very thin. In an environment where most users connected at 1,200 baud, if they were lucky, graphics and streaming video would have taken an unbearably long time to download, if they downloaded at all. What was needed was an efficient mode of communication—and text is one of the most efficient.<sup>8</sup>

Most think of this fact about the early Net as a limitation. Technically, it was. But this technical description does not exhaust its normative description as an architecture that made possible a certain kind of life. From this perspective, limitations can be features; they can enable as well as disable. And this particular limitation enabled classes of people who were disabled in real-space life.

Think about three such classes—the blind, the deaf, and the “ugly.” In real space these people face an extraordinary array of constraints on their ability to communicate. The blind person in real space is constantly confronted with architectures that presume he can see; he bears an extraordinary cost in retrofitting real-space architectures so that this presumption is not totally exclusionary. The deaf person in real space confronts architectures that presume she can hear; she too bears an extraordinary cost in retrofitting these architectures. The “ugly” person in real space (think of a bar or a social club) confronts architectures of social norms that make his appearance a barrier to a certain sort of intimacy. He endures extraordinary suffering in conforming to these architectures.

In real space these three groups are confronted with architectures that disable them relative to “the rest of us.” But in cyberspace, in its first iteration, they did not.

The blind could easily implement speech programs that read the (by definition machine-readable) text and could respond by typing. Other people on the Net would have no way of knowing that the person typing the message was blind, unless he claimed to be. The blind were equal to the seeing.

The same with the deaf. There was no need to hear anything in this early Internet. For the first time many of the deaf could have conversations, or



exchanges, in which the most salient feature was not that the person was deaf. The deaf were equal to the hearing.

And the same with the “ugly.” Because your appearance was not transmitted with every exchange, the unattractive could have an intimate conversation with others that was not automatically defined by what they looked like. They could flirt or play or be sexual without their bodies (in an extremely underappreciated sense) getting in the way. This first version of the Net made these people equal to “the beautiful.” In a virtual chat room, stunning eyes, a captivating smile, or impressive biceps don’t do it. Wit, engagement, and articulateness do.

The architecture of this original cyberspace gave these groups something that they did not have in real space. More generally, it changed the mix of benefits and burdens that people faced—the literate were enabled and the attractive disabled relative to real space. Architectures produced these enablings and disablings.

I’ve told this story as if it matters only to those who in real space are “disabled.” But of course, “disabled” is a relative term.<sup>9</sup> It is more accurate to say that the space changes the meaning of the enabled. A friend—a strikingly beautiful and powerful woman, married, and successful—described for me why she spends hours in political chat spaces, arguing with others about all sorts of political topics:

You don’t understand what it’s like to be me. You have lived your whole life in a world where your words are taken for their meaning; where what you say is heard for what it says. I’ve never had a space, before this space, where my words were taken for what they meant. Always, before, they were words of “this babe,” or “wife,” or “mother.” I could never speak as I. But here, I am as I speak.

Clearly, the space is enabling her, even though one would not have said that in real space she was “disabled.”<sup>10</sup>

Over time, as bandwidth has expanded, this architecture has changed, and so has the mix of benefits and burdens. When graphics entered the Net through the World Wide Web, the blind became “blind” again. As sound files or speech in virtual spaces have been created, the deaf have become “deaf” again. And as chat rooms have started segregating into spaces where video-cams capture real images of the people chatting and spaces where there is just text, the video-unappealing are again unappealing.<sup>11</sup> As the architectures change, definitions of who is “disabled” change as well.

My point is not to argue that the Net should not change—though of course, if it can change in ways that minimize the disabling effect of sound



and graphics, then it no doubt should.<sup>12</sup> However important, my point is not really about the “disabled” at all. I use this example simply to highlight a link—between these structures of code and the world this code enables. Codes constitute cyberspaces; spaces enable and disable individuals and groups. The selections about code are therefore in part a selection about who, what, and, most important, what ways of life will be enabled and disabled.

#### CYBER-PLACES

We can build on this point by looking at a number of “communities” that are constituted differently and that constitute different forms of life and by considering what makes these differences possible.

#### America Online

America Online (AOL) is an online service provider—“by far the largest ISP in the world”<sup>13</sup> with some 12 million subscribers in 1998 and 27 million today.<sup>14</sup> But despite having the population of New York and New Jersey combined, AOL still describes itself as a “community.” A large community perhaps, but a community nonetheless.

This community has a constitution—not in the sense of a written document (though there is that as well), but in the sense of a way of life for those who live there. Its founding vision was that community would make this place sing. So from its start, AOL’s emphasis has been on enabling people to interact, through chat, bulletin boards, and e-mail. (Today, AOL hosts the exchange of more messages daily than does the U.S. Post Office.<sup>15</sup>) Earlier providers, obsessed with providing content or advertising, limited or ignored the possibilities for interaction and exchange, but AOL saw interaction as the stuff that makes cyberspace different. It built itself on building a community and establishing itself as a place where people could say what they wanted.<sup>16</sup>

This interaction is governed by the rules of the place. Some of these rules are formal, others customary. Among the formal are express terms to which every member subscribes upon joining AOL. These terms regulate a wide range of behaviors in this space, including the behavior of AOL members anywhere on the Internet.<sup>17</sup>

Increasingly, these rules have become controversial. AOL policies have been called “Big Brother” practices. Arguments that get heated produce exchanges that are rude. But rudeness, or offensiveness, is not permitted in AOL’s community. When these exchanges are expunged, claims of “censorship” arise.<sup>18</sup>

My aim here, however, is not to criticize these rules of “netiquette.” AOL also has other rules that regulate AOL members—rules expressed not in contracts but rather through the very architectures of the space. These rules are the most important part of AOL’s constitution, but they are probably the part considered last when we think about what regulates behavior in this cyberspace.

Consider some examples:

For most of AOL’s life,<sup>19</sup> as a member of AOL you could be any one of five people. This was just one amazing feature of the space. When you started an account on AOL, you had the right to establish up to five identities, through five different “screen names” that in effect establish five different accounts. Some users, of course, used the five screen names to give other family members access to AOL. But not everyone used an AOL account like this. Think about the single woman, signing up for her first AOL account. AOL gave her up to five identities that she can define as she wishes—five different personae she can use in cyberspace.

What does that mean? A screen name is just a label for identifying who you are when you are on the system. It need not (indeed, often cannot) be your own name. If your screen name is “StrayCat,” then people can reach you by sending e-mail to “straycat@aol.com.” If you are online, people can try to talk to you by paging StrayCat on the AOL system; a dialogue would then appear on your screen asking whether you want to talk to the person who paged you. If you enter a chat room, the list of residents there will add you as “StrayCat.”

But who is StrayCat? Here is a second dimension of control. StrayCat is who StrayCat says she is. She can choose to define herself as no one at all. If she chooses to place a description of herself in the members’ directory, that description can be as complete or incomplete as she wishes. It can be true or false, explicit or vague, inviting or not. A member stumbling across StrayCat, then, in a chat room set up for stamp collectors could get her profile and read that StrayCat lives in Cleveland and is single and female. What happens next is anyone’s guess.

Yet this need only be one of StrayCat’s five identities. Let’s say there is a different persona that StrayCat likes to have when she wanders through chat rooms. She can then select another screen name and define it in the directory as she wishes. Perhaps when StrayCat is having a serious discussion in a news-group or political list she prefers to speak as herself. She could then select a screen name close to her own name and define it according to who she really is. At other times StrayCat may like to pretend to be a man—engaging in virtual cross-dressing and all that might bring with it. One of her screen names

could then be a man's. And so on. The point is the multiplicity that AOL allows, and the freedom this multiplicity permits.

No one except StrayCat needs to know which screen names are hers. She is not required to publish the full list of her identities, and no one can find out who she is (unless she breaks the rules). (After revealing to the U.S. Navy the name of one of its members so that the Navy could prosecute the person for being a homosexual, AOL adopted a very strict privacy policy that promises never to allow a similar transgression to happen again.)<sup>20</sup>

So in AOL you were given a fantastic power of pseudonymity that the "code writers" of real space simply do not give. You could, of course, try in real space to live the same range of multiple lives, and to the extent that these lives are not incompatible or inconsistent, you could quite often get away with it. For instance, you could be a Cubs fan during the summer and an opera buff during the winter. But unless you take extraordinary steps to hide your identity, in real space you are always tied back to you. You cannot simply define a different character; you must make it, and more important (and difficult), you must sustain its separation from your original identity.

That is a first feature of the constitution of AOL—a feature constituted by its code. A second is tied to speech—what you can say, and where.

Within the limits of decency, and so long as you are in the proper place, you can say what you want on AOL. But beyond these limits, speech on AOL is constrained in a more interesting way: not by rules, but by the character of the potential audience. There are places in AOL where people can gather; there are places where people can go and read messages posted by others. But there is no space where everyone gathers at one time, or even a space that everyone must sooner or later pass through. There is no public space where you could address all members of AOL. There is no town hall or town meeting where people can complain in public and have their complaints heard by others. There is no space large enough for citizens to create a riot. The owners of AOL, however, can speak to all. Steve Case, the founder of AOL, used to write "chatty" letters to the members as the community's "town mayor."<sup>21</sup> Case left AOL in 2005, and apparently no one has stepped into his speaker shoes. AOL does still advertise to all its members and can send everyone an e-mail, but only the owners and those they authorize can do so. The rest of the members of AOL can speak to crowds only where they notice a crowd—and never to a crowd greater than thirty-six (up from twenty-three when the first edition of this book was published).

This is another feature of the constitution of the space that AOL is, and it too is defined by code. That only twenty-three people can be in a chat room at once is a choice of the code engineers. While their reasons could be many,

the effect is clear. One can't imagine easily exciting members of AOL into public action, such as picketing the latest pricing policy. There are places to go to complain, but you have to take the trouble to go there yourself. There is no place where members can complain en masse.

Real space is different in this respect. Much of free speech law is devoted to preserving spaces where dissent can occur—spaces that can be noticed, and must be confronted, by nondissenting citizens.<sup>22</sup> In real space there are places where people can gather, places where they can leaflet. People have a right to the sidewalks, public streets, and other traditional public forums. They may go there and talk about issues of public import or otherwise say whatever they want. Constitutional law in real space protects the right of the passionate and the weird to get in the face of the rest. But no such design is built into AOL.<sup>23</sup> As Dawn Nunziato writes,

AOL explains in its Community Guidelines that “like any city, we take pride in—and are protective of—our community.” Unlike any other city, however, AOL enjoys the unfettered discretion to censor constitutionally-protected speech in its discussion forums and other online spaces, including “vulgar language” (which, it warns, is “no more appropriate online than [it] would be at Thanksgiving dinner”), “crude conversations about sex,” and “discussions about . . . illegal drug abuse that imply it is acceptable.”<sup>24</sup>

This is not to romanticize the power of real-space public forums. (Nor is it to pick on AOL: As Nunziato continues, “users seeking stronger protection for their expression might turn to an ISP other than AOL. They will find, however, similar restrictions on speech imposed by many other major ISPs.”<sup>25</sup>) We have become such an apolitical society that if you actually exercised this constitutionally protected right, people would think you were a nut. If you stood on a street corner and attacked the latest tax proposal in Congress, your friends would be likely to worry—and not about the tax proposal. There are exceptions—events can make salient the need for protest—but in the main, though real space has fewer controls through code on who can speak where, it has many more controls through norms on what people can say where. Perhaps in the end real space is much like AOL—the effective space for public speech is limited and often unimportant. That may well be. But my aim here is to identify the feature and to isolate what is responsible for it. And once again, it turns out to be a feature built into the code.

A third feature of AOL's constitution also comes from its code. This is traceability. While members are within the exclusive AOL content area (in other words, when they're not using AOL as a gateway to the Internet), AOL

can (and no doubt does) trace your activities and collect information about them. What files you download, what areas you frequent, who your “buddies” are—all this is available to AOL. These data are extremely valuable; they help AOL structure its space to fit customer demand. But gaining the ability to collect these data required a design decision. This decision too was part of the constitution that is AOL—again, a part constituted by its code. It is a decision that gives some but not others the power to watch.

AOL is not exclusive in this enabling capacity. It shares the power. One wonderful feature of the online space is something called “buddy lists.” Add someone to your buddy list, and when he comes online you hear the sound of a creaking door and are notified that he is online. (The “buddy” need not know he is being watched, though he can, if he knows, block the watching.) If that person goes into a chat area and you “locate” him, you will be told in what chat area he is. This power, given to ordinary users, can have complicated consequences. (Imagine sitting at work with your buddy feature turned on, watching your spouse come online, enter a chat area, and—you get the point.) This ability to monitor is built into the space. Individuals can turn it off, at least for a single watcher, but only if they know about it and think to change it.

Consider one final feature of the constitution of AOL, closely linked to the last: commerce. In AOL you can buy things. You can buy things and download them, or buy things and have them sent to your home. When you buy, you buy with a screen name, and when you buy with a screen name, AOL knows (even if no one else does) just who you are. It knows who you are, it knows where you live in real space, and most important, it knows your credit card number and the security it provides.

AOL knows who you are—this is a feature of its design. All your behavior on AOL is watched; all of it is monitored and tracked back to you as a user. AOL promises not to collect data about you individually, but it certainly collects data about you as part of a collective. And with this collective, and the link it provides back to you, AOL is a space that can better, and more efficiently, sell to you.

These four features mark AOL space as different from other places in cyberspace. It is easier for AOL to identify who you are, and harder for individuals to find out who you are; easier for AOL to speak to all its “citizens” as it wishes, and harder for dissidents to organize against AOL’s views about how things ought to be; easier for AOL to market, and harder for individuals to hide. AOL is a different normative world; it can create this different world because it is in control of the architecture of that world. Members in that space face, in a sense, a different set of laws of nature; AOL makes those laws.

Again, my aim is not to criticize the creation of this world or to say that it is improper. No doubt AOL makes promises to its members that are designed to allay some of the concern that this control creates, and no doubt if the place became oppressive, the market would provide plenty of alternatives.

Rather my objective is to impart a sense of what makes AOL the way it is. It is not just written rules; it is not just custom; it is not just the supply and demand of a knowing consuming public. What makes AOL is in large part the structure of the space. You enter AOL and you find it to be a certain universe. This space is constituted by its code. You can resist this code—you can resist how you find it, just as you can resist cold weather by putting on a sweater. But you are not going to change how it is. You do not have the power to change AOL's code, and there is no place where you could rally AOL members to force AOL to change the code. You live life in AOL subject to its terms; if you do not like them, you go elsewhere.

These features of the AOL space have important implications for how it is regulated. Imagine there is a problem on AOL that AOL wants to stop. It wants to prevent or at least control a certain behavior. What tools does AOL have?

First, it has all the tools that any club, fraternity, or "community" might have. It can announce rules for its members (and AOL certainly does). Or it can try to stigmatize the behavior, to use the norms of the community to help regulate the problem. This AOL does as well. Alternatively, if the problem comes from the overuse of a particular resource, then the managers at AOL can price that resource differently by exacting a tax to reduce its usage or a different price for those who use it too much.

But AOL has something more at hand. If AOL does not like a certain behavior, then in at least some cases it can regulate that behavior by changing its architecture. If AOL is trying to control indecent language, it can write routines that monitor language usage; if there is improper mixing between adults and kids, AOL can track who is talking to whom; if there is a virus problem caused by people uploading infected files, it can run the files automatically through virus checkers; if there is stalking or harassing or threatening behavior, AOL can block the connection between any two individuals.

In short, AOL can deal with certain types of problems by changing its code. Because the universe that AOL members know (while in AOL) is defined by this code, AOL can use the code to regulate its members.

Think a bit about the power I am describing—and again, I am not complaining or criticizing or questioning this power, only describing it. As you move through this space that AOL defines—entering a chat area, posting a

message to a bulletin board, entering a discussion space, sending instant-messages to another person, watching or following other people, uploading or downloading files from sites, turning to certain channels and reading certain articles, or obsessively paging through a space looking for pictures of a certain actor or actress—as you do any of these things, AOL is, in an important sense, there. It is as if the system gives you a space suit that you use to navigate the space but that simultaneously monitors your every move.

In principle, the potential for control is extraordinary. Imagine AOL slowing the response time for a certain kind of service it wants to discourage, or channeling the surfer through ads that it wants customers to see, or identifying patterns of behavior that its monitors would watch, based on the fear that people with patterns like X are typically dangerous to people of type Y. I do not think AOL engages in activities like these, and I am not saying that there would be anything wrong if it did. But it is important to note that the potential for control in this “community” is unlimited—not in the sense that AOL could make life miserable (since people would then leave), but in the sense that it has a regulatory tool that others, in both real space and other cyberspaces, do not. Its power is, of course, checked by the market, but it has a tool of control that others in the market, but outside cyberspace, do not have.

In principle, then, AOL must choose. Every time AOL decides that it wants to regulate a certain kind of behavior, it must select from among at least four modalities—rules, norms, prices, or architecture. And when selecting one of these four modalities, selecting architecture as a regulator will often make the most sense.

#### Counsel Connect

David Johnson began Counsel Connect (CC) in 1992 as an online lawyers’ cooperative. The idea was simple: Give subscribers access to each other; let them engage in conversations with each other; and through this access and these conversations, value would be created. Lawyers would give and take work; they would contribute ideas as they found ideas in the space. A different kind of law practice would emerge—less insular, less exclusive, more broadly based.

I thought the idea amazing, though many thought it nuts. For a time the system was carried by Lexis; in 1996 it was sold to American Lawyer Media, L.P.; in 1997 it migrated to the Internet, and it closed in 1999.<sup>26</sup> At its peak, it boasted thousands of subscribers, though it is hard to know how many of them contributed to the discussion online. Most simply watched the discus-

sions of others, perhaps linking three or four discussion groups of their particular interest, plus a few of more general interest. But many saw the emerging culture as something amazing and new (for lawyers at least). As its founder, David Johnson, described it, “Think of The Well for lawyers, with its own highly unique evolution, emergence, maintenance, and adaptation.”<sup>27</sup> Members got to know each other well. “Inevitably, this led to numerous real world meetings. . . . Of those I attended, it always resemble[d] a get together of long-time acquaintances even though many of us ha[d] not previously met face to face.”<sup>28</sup>

The discussion was organized into legal topics. Each topic was divided into discussion groups, with each group led by a discussion leader. The leader was not a moderator; he or she had no power to cancel a post. The leader was there to inspire conversation—to induce others to speak by being encouraging or provocative.

At its height, there were some 90 groups in this space. The poster of a particular message may have had it removed, but if the poster did not remove it, it stayed—at first in the list of topics being discussed, and later in an archive that could be searched by any member.

Members paid a fee to join and get an account with their real name on it. Postings use members’ real names, and anyone wondering who someone is could simply link to a directory. Members of CC must be members of the bar, unless they are journalists. Others have no right to access; the community here is exclusive.

Postings in the space look very much like postings in a USENET newsgroup. A thread could be started by anyone, and replies to a thread were appended to the end. Because messages did not move off the system, one could easily read from the start of a thread to its end. The whole conversation, not just a snippet, was preserved.

These features of CC space were obviously designed; the architects chose to enable certain features and to disable others. We can list here some of the effects of these choices.

First, there was the effect from being required to use your own name. You were more likely to think before speaking and to be careful about being right before saying something definitive. You were constrained by the community, which would judge what you said, and in this community you could not escape from being linked to what you said. Responsibility was a consequence of this architecture, but so was a certain inhibition. Does a senior partner at a leading law firm really want to ask a question that will announce his ignorance about a certain area of law? Names cannot be changed to protect the ignorant, so they will often simply not speak.

Second, there was an effect from forcing all discussion into threads. Postings were kept together; a question was asked, and the discussion began from the question. If you wanted to contribute to this discussion, you had to first read through the other postings before responding. Of course, this was not a technical requirement—you certainly had a choice not to read. But if you did not read through the entire thread, you could well end up repeating what another had said and so reveal that you were speaking without listening. Again, the use of real names ties members' behavior to the norms of the community.

Third, there was the effect of reputation: The reputation you built in this space was based on the kind of advice you gave. Your reputation survived any particular post and was, of course, affected by any subsequent posts. These posts were archived and searchable. If you said one thing about topic X and then the opposite later on, you were at least open to a question about consistency.

Fourth, there was the effect of tying reputation to a real name in a real community of professionals. Misbehaving here mattered elsewhere. CC thus got the benefit of that community—it got the benefit, that is, of the norms of a particular community. These norms might have supported relatively productive community behavior—more productive, that is, than the behavior of a group whose members are fundamentally mixed. They might also have supported punishing those who deviated from appropriate behavior. Thus, CC got the benefit of community sanction to control improper behavior, whereas AOL had to rely on its own content police to ensure that people stayed properly on topic.

We can describe the world of CC that these features constitute in two different ways, just as we can describe the world AOL constitutes in two different ways. One is the life that CC's features made possible—highly dialogic and engaged, but monitored and with consequences. The other is the regulability by the manager of the life that goes on in the CC space. And here we can see a significant difference between this space and AOL.

CC could have used the norms of a community to regulate more effectively than AOL can. CC benefited from the norms of the legal community; it knew that any misbehavior would be sanctioned by that community. There was, of course, less “behavior” in this space than in AOL (you did fewer things here), but such as it was, CC behavior was quite significantly regulated by the reputations of members and the consequences of using their real names.

These differences together had an effect on CC's ability to regulate its members. They enabled a regulation through modalities other than code.



They made behavior in CC more regulable by norms than behavior in AOL is. CC in turn may have had less control than AOL does (since the controlling norms are those of the legal community), but it also bore less of the burden of regulating its members' behavior. Limiting the population, making members' behavior public, tying them to their real names—these are the tools of self-regulation in this virtual space.

But CC was like AOL in one important way: It was not a democracy and neither is AOL. Management in both cases controls what will happen in the space—again, not without constraint, because the market is an important constraint. But in neither place do “the people” have the power to control what goes on. Perhaps they did, indirectly, in CC more than AOL, since it is the norms of “the people” that regulate behavior in CC. But these norms cannot be used against CC directly. The decisions of CC and AOL managers may have been affected by market forces—individuals can exit, competitors can steal customers away. But voting doesn't direct where AOL goes, and it didn't with CC either.

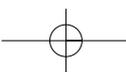
That's not the case with the next cyber-place. At least, not anymore.

#### LambdaMOO

LambdaMOO is a text-based virtual reality. People from across the world (today close to six thousand of them) link to this space and interact in ways that the space permits. The reality is the product of this interaction. Individuals can participate in the construction of this reality—sometimes for upwards of eighty hours a week. For some this interaction is the most sustained human contact of their entire lives. For most it is a kind of interaction unmatched by anything else they know.

In the main, people just talk here. But it is not the talk of an AOL chat room. The talk in a MUD is in the service of construction—of constructing a character and a community. You interact in part by talking, and this talking is tied to a name. This name, and the memories of what it has done, live in the space, and over time people in the space come to know the person by what these memories recall.

The life within these MUDs differ. Elizabeth Reid describes two different “styles”<sup>29</sup>—social-style MUD and an adventure or game-style MUD. Social MUDs are simply online communities where people talk and build characters or elements for the MUD. Adventure MUDs are games, with (virtual) prizes or power to be won through the deployment of skill in capturing resources or defeating an enemy. In either context, the communities survive a particular interaction. They become virtual clubs, though with different purposes. Members build reputations through their behavior in these clubs.



You get a character simply by joining the MOO (though in LambdaMOO the waiting list for a character extends over many months). When you join the space, you define the character you will have. At least, you define certain features of your character. You select a name and a gender (no gender is an option as well) and describe your character. Some descriptions are quite ordinary (Johnny Manhattan is “tall and thin, pale as string cheese, wearing a neighborhood hat”).<sup>30</sup> Others, however, are quite extraordinary. (Legba, for instance, is a Haitian trickster spirit of indeterminate gender, brown-skinned and wearing an expensive pearl gray suit, top hat, and dark glasses.)<sup>31</sup>

Julian Dibbell broke the story of this space to the nonvirtual world in an article in the *Village Voice*.<sup>32</sup> The story that was the focus of Dibbell’s article involved a character called Mr. Bungle who, it turns out, was actually a group of NYU undergraduates sharing this single identity. Bungle entered a room late one evening and found a group of characters well known in that space. The full story cannot be told any better than Dibbell tells it. For our purposes, the facts will be enough.<sup>33</sup>

Bungle had a special sort of power. By earning special standing in the LambdaMOO community, he had “voodoo” power: he could take over the voices and actions of other characters and make them appear to do things they did not really do. This Bungle did that night to a group of women and at least one person of ambiguous gender. He invoked this power, in this public space, and took over the voices of these people. Once they were in his control, Bungle “raped” these women, violently and sadistically, and made it seem as if they enjoyed the rape.

The “rape” was virtual in the sense that the event happened only on the wires. “No bodies touched,” as Dibbell describes it.

Whatever physical interaction occurred consisted of a mingling of electronic signals sent from sites spread out between New York City and Sydney, Australia. . . . He commenced his assault entirely unprovoked at, or about 10 P.M. Pacific Standard Time. . . . [H]e began by using his voodoo doll to force one of the room’s occupants to sexually service him in a variety of more or less conventional ways. That this victim was exu. . . . He turned his attentions now to Moon-dreamer . . . forcing her into unwanted liaisons with other individuals present in the room. . . . His actions grew progressively violent. . . . He caused Moon-dreamer to violate herself with a piece of kitchen cutlery. He could not be stopped until at last someone summoned Iggy . . . who brought with him a gun of near wizardly powers, a gun that didn’t kill but enveloped its targets in a cage impermeable even to a voodoo doll’s powers.<sup>34</sup>



Rape is a difficult word to use in any context, but particularly here. Some will object that whatever happened in this virtual space, it has nothing to do with rape. Yet even if “it” was not “rape,” all will see a link between rape and what happened to these women there. Bungle used his power over these women for his own (and against their) sexual desire; he sexualized his violence and denied them even the dignity of registering their protest.

For our purposes, whether what happened here was really rape is beside the point. What matters is how the community reacted. The community was outraged by what Bungle had done, and many thought something should be done in response.

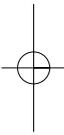
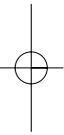
They gathered, this community of members of LambdaMOO, in a virtual room at a set time, to discuss what to do. Some thirty showed up, the largest meeting the community had known. Some thought that Bungle should be expelled—“toaded,” as it is described, killed for purposes of the MOO. Others thought that nothing should be done; Bungle was certainly a creep, but the best thing to do to creeps was simply to ignore them. Some called on the Wizards of the space—the creators, the gods—to intervene to deal with this character. The Wizards declined: Their job, they replied, was to create the world; the members had to learn to live within it.

There was really no law that governed what Bungle had done. No real-space law reached sexual pranks like this, and neither did any explicit rule of LambdaMOO.<sup>35</sup> This troubled many who wanted to do something. Invoking real-space ideals about fair notice and due process, these people argued that Bungle could not be punished for violating rules that did not exist at the time.

Two extremes eventually emerged. One side urged vigilantism: Bungle was a miscreant, and something should be done about him. But what shouldn’t be done, they argued, was for LambdaMOO to respond by creating a world of regulation. LambdaMOO did not need a state; it needed a few good vigilantes. It needed people who would enforce the will of the community without the permanent intrusion of some central force called the state. Bungle should be expelled, killed, or “toaded”—and someone would do it. But only if the group resisted the call to organize itself into a state.

The other side promoted just one idea: democracy. With the cooperation of the Wizards, LambdaMOO should establish a way to vote on rules that would govern how people in the space behaved. Any question could be made the subject of a ballot; there was no constitution limiting the scope of what democracy could decide. An issue decided by the ballot would be implemented by the Wizards. From then on, it would be a rule.

Both extremes had their virtues, and both invited certain vices. The anarchy of the first risked chaos. It was easy to imagine the community turning



against people with little or no warning; one imagined vigilantes roaming the space, unconstrained by any rules, “toading” people whose crimes happened to strike them as “awful.” For those who took this place less seriously than real space, this compromise was tolerable. But what was tolerable for some was intolerable to others—as Bungle had learned.

Democracy seemed natural, yet many resisted it as well. The idea that politics could exist in LambdaMOO seemed to sully the space. The thought that ideas would have to be debated and then voted on was just another burden. Sure, rules would be known and behavior could be regulated, but it all began to seem like work. The work took something from the fun the space was to have been.

In the end, both happened. The debate that evening wound down after almost three hours. No clear resolution had found its way in. But a resolution of sorts did occur. As Dibbell describes it:

It was also at this point, most likely, that TomTraceback reached his decision. TomTraceback was a wizard, a taciturn sort of fellow who’d sat brooding on the sidelines all evening. He hadn’t said a lot, but what he had said indicated that he took the crime committed against exu and Moondreamer very seriously, and that he felt no particular compassion toward the character who had committed it. But on the other hand he had made it equally plain that he took the elimination of a fellow player just as seriously, and moreover that he had no desire to return to the days of wizardly intervention. It must have been difficult, therefore, to reconcile the conflicting impulses churning within him at that moment. In fact, it was probably impossible, for . . . as much as he would have liked to make himself an instrument of the MOO’s collective will, [he surely realized that under the present order of things] he must in the final analysis either act alone or not act at all.

So TomTraceback acted alone.

He told the lingering few players in the room that he had to go, and then he went. It was a minute or two before 10 P.M. He did it quietly and he did it privately, but all anyone had to do to know he’d done it was to type the @who command, which was normally what you typed if you wanted to know a player’s present location and the time he last logged in. But if you had run an @who on Mr. Bungle not too long after TomTraceback left emmeline’s room, the database would have told you something different.

“Mr\_Bungle,” it would have said, “is not the name of any player.”

The date, as it happened, was April Fool’s Day, but this was no joke: Mr. Bungle was truly dead and truly gone.<sup>36</sup>

When the Wizards saw this, they moved to the other extreme. With no formal decision by the citizens, the Wizards called forth a democracy. Starting May 1, 1993,<sup>37</sup> any matter could be decided by ballot, and any proposition receiving at least twice as many votes for as against would become the law.<sup>38</sup> Many wondered whether this was an advance or not.

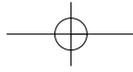
There is a lot to think about in this story, even in my savagely abridged version.<sup>39</sup> But I want to focus on the sense of loss that accompanied the Wizards' decision. There is a certain romance tied to the idea of establishing a democracy—Kodak commercials with tearful Berliners as the Wall comes down and all that. The romance is the idea of self-government and of establishing structures that facilitate it. But LambdaMOO's move to self-government, through structures of democracy, was not just an achievement. It was also a failure. The space had failed. It had failed, we could say, to self-regulate. It had failed to engender values in its population sufficient to avoid just the sort of evil Bungle had perpetrated. The debate marked the passage of the space from one kind of place to another. From a space self-regulated to a space regulated by self.

It might seem odd that there would be a place where the emergence of democracy would so depress people. But this kind of reaction is not uncommon in cyber-places. Katie Hafner and Matthew Lyon tell a story of the emergence of a "widget" called the FINGER command on UNIX, that would allow users to see when the last time another user had been on the computer, and whether she had read her mail. Some thought (not surprisingly, I should think) that this command was something of an invasion of privacy. Whose business was it when I was last at my machine, and why should they get to know whether I have read my mail?

A programmer at Carnegie Mellon University, Ivor Durham, changed the command to give the user the power to avoid this spying finger. The result? "Durham was flamed without mercy. He was called everything from spineless to socially irresponsible to a petty politician, and worse—but not for protecting privacy. He was criticized for monkeying with the openness of the network."<sup>40</sup>

The values of the UNIX world were different. They were values embedded in the code of UNIX. To change the code was to change the values, and members of the community fought that change.

So too with the changes to LambdaMOO. Before the balloting, LambdaMOO was regulated through norms. These regulations of social structures were sustained by the constant policing of individual citizens. They were the regulations of a community; the rise of democracy marked the fall of this community. Although norms would no doubt survive the establishment of a



democracy, their status was forever changed. Before the democracy, a struggle over which norms should prevail could be resolved only by consensus—by certain views prevailing in a decentralized way. Now such a struggle could be resolved by the power of a majority—not through what a majority did, but through how they voted.

I've romanticized this bizarre little world far more than I intended. I do not mean to suggest that the world of LambdaMOO before democracy was necessarily better than the one after. I want only to mark a particular change. Like CC, and unlike AOL, LambdaMOO is a place where norms regulate. But unlike CC, LambdaMOO is now a place where members have control over restructuring the norms.

Such control changes things. Norms become different when ballots can overrule them, and code becomes different when ballots can order Wizards to change the world. These changes mark a movement from one kind of normative space to another, from one kind of regulation to another.

In all three of these cyber-places, code is a regulator. But there are important differences among the three. Norms have a relevance in CC and LambdaMOO that they do not in AOL; democracy has a relevance in LambdaMOO that it does not have in CC or AOL. And monitoring has a relevance in AOL that it does not have in LambdaMOO or CC (since neither of the latter two use data about individuals for commercial purposes, either internal or external to the organization). Code constitutes these three communities; as Jennifer Mnookin says of LambdaMOO, "politics [is] implemented through technology."<sup>41</sup> Differences in the code constitute them differently, but some code makes community thicker than others. Where community is thick, norms can regulate.

The next space in this survey is also constituted by code, though in this case the "management" has less ability to change its basic architecture. This code is net code—a protocol of the Internet that is not easily changed by a single user. At least it was not easy for me.

.law.cyber

His name was IBEX, and no one knew who he was. I probably could have figured it out—I had the data to track him down—but after he did what he did, I did not want to know who he was. He was probably a student in the very first class about cyberspace that I taught, and I would have failed him, because I was furious about what he had done. The class was "The Law of Cyberspace"; version one of that class was at Yale.

I say version one because I had the extraordinary opportunity to teach that class at three extraordinary law schools—first at Yale, then at the University of



Chicago, and finally at Harvard. These were three very different places, with three very different student bodies, but one part of the course was the same in each place. Every year a “newsgroup” was associated with the class—an electronic bulletin board where students could post messages about questions raised in the course, or about anything at all. These postings began conversations—threads of discussion, one message posted after another, debating or questioning what the earlier message had said.

These newsgroups constituted what philosophers might call “dialogic communities.” They were spaces where discussion could occur, but where what was said was preserved for others to read, as in CC. That was the dialogic part. The community was what was made over time as people got to know each other—both in this space and in real space. One year students in the class and students outside the class (who had been watching the .law.cyber discussions develop) had a party; another year the students outside the class were invited to attend one class. But over the three years, at three different schools, it was clear that three communities had been made. Each was born on a particular date, and each lived for at least a couple of months.

My story here comes from Yale. Yale is an odd sort of law school, though odd in a good way. It is small and filled with extremely bright people, many of whom do not really want to be lawyers. It fashions itself as a community, and everyone from the dean on down (not a “Yale” way to describe things) strives continuously to foster and sustain this sense of community among the students. To a large extent, it works—not in the sense that there is perpetual peace, but in the sense that people everywhere are aware of this sense of community. Some embrace it, others resist it, but resistance, like an embrace, says that something is there. One does not resist the community of people on a Greyhound bus.

One extraordinary feature of the Yale Law School is “the Wall.” The Wall is a place where people can post comments about whatever they want to say. A letter can be posted about gay rights at Yale, or a protest about Yale’s treatment of unionized workers. Political messages are posted as well as points about law. Each posting makes additional ones possible—either scribbled on the original post or appended underneath the post.

An extraordinary sign for any visitor, the Wall is located right at the center of the law school. In the middle of a fake Gothic structure is a stone space with scores of papers posted in random fashion. Around the posts stand wandering students, reading what others have said. This is Yale’s speakers’ corner, though the speakers are writers, and the writing is substantive. There is little to be gained on the Wall through rhetoric; to gain respect there, you must say something of substance.

One rule, however, governs this space. All postings must be signed; any posting without a signature is removed. Originally, no doubt, the rule meant that the posting must be signed by the person who wrote it. But because this is Yale, where no rule can exist without a thousand questions raised, a custom has emerged whereby an anonymous post can be signed by someone not its author (“Signed but not written by X”). That signature gives the post the pedigree it needs to survive on the Wall.

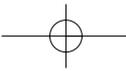
The reasons for this rule are clear, but so too are its problems. Let’s say you want to criticize the dean for a decision he has made. The dean, however sweet, is a powerful person, and you might well prefer to post a message without your name attached to it. Or say you are a student with political views that make you an outsider. Posting a message with those views and your signature might draw the scorn of your classmates. Free speech is not speech without consequence, and scorn, or shame, or ostracism are likely consequences of lots of speech.

Anonymity, then, is a way around this dilemma. With anonymity, you can say what you want without fear. In some cases, for some people, the right to speak anonymously makes sense.

Still, a community might want to resist this right. Just as anonymity might give you the strength to state an unpopular view, it can also shield you if you post an irresponsible, or slanderous, or hurtful view. You might want to question the policies of the dean, or you might want falsely to accuse a fellow student of cheating. Both utterances benefit from anonymity, but the community has good reason to resist utterances like the second.

As far as I know, IBEX never said anything on the Wall. Instead, he spoke in the newsgroup associated with my class. By design, the newsgroup was open to anyone at Yale who wanted to speak. Unlike the Wall, however, the technology allowed users to call themselves whatever they wanted. “IBEX,” of course, was a pseudonym. For purposes of the Wall, a pseudonym was just like anonymous speech—you did not have to use your real name. But in a newsgroup a pseudonymous posting is quite different from an anonymous posting. Over time you can come to know the character of a pseudonym. In the class that year, along with IBEX, we had SpeedRacer, MadMacs, CliffClaven, Aliens, blah, and Christopher Robbin. While members of the class might know who these participants were (we all knew who MadMacs was, but only a few of us knew SpeedRacer), each pseudonym had a character.

The character of IBEX was bad; this much was clear from the start. Before IBEX appeared, life in the space flourished. At first people were timid, but polite. Brave souls would post an idea or a joke, and conversation would continue around the idea or joke for a bit. After a couple of weeks the conversa-



tion would become quite intense. Patterns of exchange began. People had questions; others had answers. People stumbled as they spoke, but they were beginning, slowly, to speak.

Some things about how they spoke were immediately noticeable. First, women spoke more in this space than they did in class. Maybe not more in a statistically significant sense, but more.<sup>42</sup> Second, helpers quickly developed and differentiated from those who received their help. Soon a class developed online—a real class that identified itself as such and spoke as a class in a way that a teacher dreams of in real space, and in a way I had never known.

Why this happened I could not really say. Una Smith may have been a catalyst. I said that I taught this course three times. Each time (without my intervention at all) there was an Una Smith participating in the newsgroup. At Yale she was a real person, but after Yale I thought of her as a type. She was always a woman from outside the class; she was always extremely knowledgeable about the Net and about USENET; and she always wandered into my (virtual) class and began telling the others how they should behave. When someone violated a norm of the Net, Una would correct them. Often this instruction was not taken terribly well (these were, after all, law students). Soon the class would rally to defend the instructed and to challenge her to defend her rules. And of course, expert that she was, she usually had an answer that did defend the rules she had dictated. This exchange soon became a focus of the class. Una had drawn their anger, and the class gained cohesiveness as a result.

About a month and a half into the course, the group reached an apex of sorts. It became the best it would be. I remember the moment well. Early on a spring afternoon I noticed that someone had posted the first line of a poem. By the end of the day, without any coordination, the class had finished the poem. There had been rhythm to the exchanges; now there was rhyme. Things hummed in the newsgroup, and people were genuinely surprised about this space.

It was then that IBEX appeared. I think it was just after we had discussed anonymity in class, so maybe his later claims to have been serving a pedagogical role were true. But he appeared after one of our classes—appeared, it seemed, just to issue an attack on another member of the class. Not an attack on his ideas, but on him. So vicious and so extensive was this attack that when I read it, I didn't know quite how to understand it. Could it have been real?

Almost immediately, conversation in the group died. It just stopped. No one said anything, as if everyone were afraid that the monster that had entered our space would turn his fury on one of them next. Until, that is, the victim responded, with an answer that evinced the wounds of the attack. IBEX's words had cut. The victim was angry and hurt, and he attacked back.



But his salvo only inspired another round of viciousness, even more vile than the first. With this, other members of the class could not resist joining in. IBEX was attacked by a string of characters in the class as cowardly for hiding behind a pseudonym and as sick for what he had said. None of this had any effect. IBEX came back, again and again, with an ugliness that was as extreme as it was unrelenting.

The space had been changed. Conversation fell off, people drifted away. Some no doubt left because they were disgusted with what had happened; others did not want to be IBEX's next target. There was a brief period of life in the space as people rallied to attack IBEX. But as he came back again and again, each time more vicious than the last, most simply left. (One time IBEX came back to protest that he had been wronged; in the week before, he claimed, he had not posted anything, but someone wearing the white sheet of IBEX had posted in IBEX's name, so that he, the real IBEX, had been defamed. The class had little sympathy.)

But it was not just the online class that changed. As we met face to face each week, I felt the atmosphere bend. People felt the creature in the room, though no one could believe he was a student at the Yale Law School. This was their classmate, hiding behind a smile or a joke in real space, but vicious in cyberspace. And the very idea that this evil was hidden under a smile changed how people felt about smiles.

Some called this the "David Lynch effect," an allusion to the director who portrays the rot of society just under freshly painted façades. We felt in that class the rot of our community just under the surface of smiling and functional students. There was a (relatively tame) Jake Baker in our midst. The space had permitted behavior that destroyed community—community that the space itself had created. Community had been created in part through the ability to hide—to hide behind a benign pseudonym; to hide hesitation, or editing, in the writing; to hide your reaction; to hide that you were not paying attention. These anonymities had made the community what it was. But the same anonymity that created the community gave birth to IBEX as well, and thus took the community away.

SecondLi(f/v)e(s)

These four places that I have just described were all described in the first edition of this book, each in just about the same terms. They're old stories, and the lessons they teach are still precisely the lesson this chapter is meant to convey. But I don't mean to suggest that there's been no interesting progress in the cyberspaces that the Internet has inspired. The last five years have wit-



nessed an explosion in cyberspaces, much more dramatic than anything I imagined when I first wrote this book.

In one sense, these spaces are nothing really new. They have fancy new technology that, because computers are faster and bandwidth is broader, functions much better than their earlier versions. But the MMOG space I described in Chapter 2 was inspired by real places.

What's changed, however, is size. As Julian Dibbell described it to me, the question is

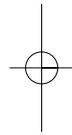
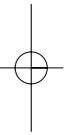
does size matter in these kinds of spaces? And I think it does. The text-based world is naturally limited in size. The limit is not so much text versus graphics as it is limited cultural accessibility versus a much broader accessibility. That makes for larger spaces.<sup>43</sup>

The result is “something socially richer in a lot of ways,” “not so much the particular affordances of 3D graphic imagery, which will also someday look pretty crude.”

Massively Multiple Online Role Playing Games (again, MMOGs, or MMORPGs) have become a whole industry. Literally millions spend hundreds, sometimes thousands of hours each year in these spaces along with literally billions of dollars to live these second lives. While living these second lives, of course, they are also living a life in real space. When they're playing the MMOG World of Warcraft, they are at the same time playing father or wife in real space. They have thus not left the real world to go to these other places. But they integrate the other places into their real world life, and the last five years has seen an explosion in the percentage of real-world life that is lived virtually.

These “games” can be divided roughly into two types. In one type, people “play” a game that has been defined by others. These are “role-playing games.” Thus, World of Warcraft is a role-playing game in which people compete to gain wealth and status (making it not so different from real life). Grand Theft Auto is a game in which people engage in a kind of virtual crime. These games all have a structure to them, but they differ in the degree to which people can customize or create their own characters or environments. The vast majority of online games are role-playing games in this sense. One site that tracks these communities estimates 97 percent are role-playing games of some sort.<sup>44</sup>

The second type involves much more construction. These spaces provide communities in which people at a minimum socialize. In addition to socializing, there is creative and commercial activity. Depending upon the game, the mix among these activities differs substantially. But they all aim to create



a virtual world that inspires a real community within itself. These games are an extension of the MOOs I described above. But they extend the virtual community of a MOO beyond those who feel comfortable manipulating text. These worlds are graphically real, even if they are virtual.

Of course, within both of these types of MMOGs, there is creativity. The differences between them are simply a matter of degree. And within both, there is commerce. Second Life—described more below—generates over “\$4,000,000 U.S. in interpersonal transactions”<sup>45</sup> a month. Aggregated across games, as Edward Castronova describes, there is a great deal of commerce produced by these virtual worlds.

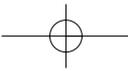
“The commerce flow generated by people buying and selling money and other virtual items (that is, magic wands, spaceships, armor) amounts to at least \$30 million annually in the United States, and \$100 million globally.”<sup>46</sup>

And more interesting (and bizarre) is Castronova’s estimate of the gross national product per capita produced in various virtual worlds. EverQuest, for example, has a GDP which is about half that of “the Caribbean Island Nation of Dominica.”<sup>47</sup> And the GDP per capita of Norrath “was about the same as Bulgaria’s and four times higher than China’s or India’s.”<sup>48</sup>

For my purposes here, however, I want to focus on the second type of MMOG, and two of these in particular. The first was an early leader in this space—There. The second is a growing and extraordinary success—Second Life.

Second Life is, as its website describes, “a 3-D virtual world entirely built and owned by its residents.” 3-D in the sense that the experience seems three dimensional—the characters and the objects appear to be in three dimensions. A *virtual* world in the sense that the objects and people are rendered by computers. *Built* by its residents in the sense that Second Life merely provided a platform upon which its residents built the Second Life world. (And not just a few. On any given day, 15 percent of Second Life residents are editing the scripts that make Second Life run.<sup>49</sup> That platform originally rendered beautiful green fields. Residents acquired land in that world, and began building structures.) And *owned* by its residents in the sense that the stuff that the residents of Second Life build is theirs—both the “physical” thing itself (the car, or the surfboard, or the house), and any intellectual property right which might be embedded in that thing that they have built.

It is this last feature that contrasts most interestingly (for me at least) with the other MMOG that I mentioned, There. There was also a community site. But it was a radically different (and less successful) world from Second Life. It was to be centered around corporate franchises—Sony or Nike, for



example, were expected to set up shop in There. People would also be allowed to create things in There, and when they sold or gave them away, There would get a percentage. The space itself came much more pre-fab, but there was significant opportunity for customization.

Its founders crafted the rhetoric of There at least around (at least their understanding of) the ideals of the United States. The exchange rate for There-bucks was 1787 to 1—1787 being the year the United States Constitution was written. And as the then-CEO of There explained to a class I was teaching, the values of the American republic informed the values of There.

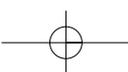
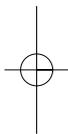
My students were skeptical. And one fantastically bright student, Catherine Crump, gave the CEO a bit of a rough ride. She asked whether There would respect the principles of the First Amendment. “Of course,” responded the CEO. “Would a citizen of There be allowed to put a sign on his land?” “Of course.” “Would she be allowed to buy land next to, say, Nike?” “Of course.” “Would she be allowed to put a sign up on her land next to Nike that says ‘Nike uses sweatshop labor?’” “Umm. I’m not sure about that.” So much for the First Amendment.

Or more relevantly to Second Life, Crump asked, “Who owns the IP [intellectual property] in the designs a citizen creates?” “There does.” “Who owns the IP in the designs Nike creates?” “Of course, Nike does. How could it be any other way?” Well, it could be another way if you followed the principles of the American Constitution, Crump suggested, which said IP rights get vested in “authors or inventors,” not in corporations.

There’s real problem, however, was structural. It is the same problem of any planned or centralized economy. There was to be built by There, Inc. And therein was its problem. The structures of these virtual worlds are extraordinarily complex. The cost of building them is immense, and thus There, Inc. faced a huge capital cost in making There run.

Second Life (like all new nations) outsourced that cost of construction to its citizens. When you buy land in Second Life, you get an empty field or deserted island. You then have to buy, barter, or build to make it habitable. There’s an economy to building it, and it can be hard work. But the things you build you can sell. And again, the designs you make are yours. More than 100,000 people now inhabit, and construct, Second Life. For them, the game is what it says.

These current rules, however, are the product of an evolution in Second Life. In the first public Alpha testing of the site that would become Second Life, there was no concept of land ownership. Everything was public. The ownership of land began with Beta testing, when all users could claim the public land at a price. When the land was claimed, the user could select



whether others could create objects, scripts, or landmarks for the land. Later the options were extended.

In version 1.1, there was a fairly major change to the physics of land. Whereas before users were free to teleport anywhere, now, to avoid harassment, owners of land could decide whether others could “trespass” or not—either by setting a default to grant or deny access, or by adding a list of people who were free to visit. These restrictions, however, applied only to the first 15 meters above the property. Beyond that, anyone was free to fly, even if the owner didn’t want them on the property.

Now this last restriction has an interesting parallel to the history of American law. As I describe in *Free Culture*,<sup>50</sup> property law in the American tradition considered the owner of land the owner of the space from the ground “an indefinite extent, upwards.”<sup>51</sup> This created an obvious conflict when airplanes appeared. Did the pilot of an airplane trespass when he flew over your land?

The accommodation the law eventually drew was between flying very low and flying very high. It was not trespassing to fly very high over someone’s land; it was a nuisance to fly very low over someone’s land. So something like the solution that Second Life achieved was also achieved by the law.

But notice the important difference. In real space, the law means you can be penalized for violating the “high/low” rule. In Second Life, you simply can’t violate the 15-meter rule. The rule is part of the code. The code controls how you are in Second Life. There isn’t a choice about obeying the rule or not, any more than there’s a choice about obeying gravity.

So code is law here. That code/law enforces its control directly. But obviously, this code (like law) changes. The key is to recognize that this change in the code is (unlike the laws of nature) crafted to reflect choices and values of the coders.

Consider another illustration of the same point. As I said, Second Life gives the creators of Intellectual Property in Second Life ownership of that property—both inside and outside Second Life.<sup>52</sup> (As one of the founders described, “Our lawyers shook their heads, but we decided the future of our company isn’t tied up in our owning what our users create.”<sup>53</sup>) That’s the same with IP in real space: Unless you’ve signed your rights away to a corporation (don’t!), when you create in real space, the law automatically gives you a copyright in your creativity. In both spaces, too, you have the right to give those rights away. I run a nonprofit called Creative Commons that makes it simple for creators to signal the freedoms they want to run with their creativity. In real space, when you use a Creative Commons license, you mark your content with the license you want. Users then know the freedoms they have. If a right is violated, it gets remedied through the law.



Second Life has taken this idea one step further. Creators in Second Life can mark their content with the license they want. But the wizards of this world are exploring the idea that the license they've selected could affect directly what others can do with that creativity. If content is marked with a Creative Commons license, then someone can take a picture of it without express permission. But if it is not marked with a license, then if you try to take a picture of it, the object will be invisible. Here again, the code expresses the law more effectively than the law in real space ever could.

#### The Internet

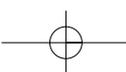
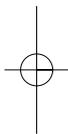
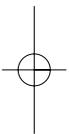
As I said, we can distinguish cyberspace from the Internet. But the point of this chapter, however clear with respect to cyberspace, is still true of the Internet. There are architectural features of the Internet that embed certain values. Those features can also change, and if they do, the values the Internet promotes will be different.

The most significant example of this is one I only mentioned in the first edition of this book, but which was at the center of *The Future of Ideas*. This is the “end-to-end” principle described by network architects Jerome Saltzer, David Clark, and David Reed in 1981.<sup>54</sup> The end-to-end (“e2e”) principle is a design philosophy about how networks should be built. It counsels that a network should be kept as simple as possible and that the intelligence required in a network be vested in the edge, or ends of a network, at least so far as that's possible.

As I've already described, the Internet embodied this principle by keeping the functionality of TCP/IP focused quite narrowly—that is, on the single function best-efforts delivery of packets of data. What those packets do, or who they're meant for, is not a concern of the protocol. Just delivering packets is the end.

One consequence of this design, then, is that people can innovate for this network without any need to coordinate with any network owner. If you want to develop an application to deliver voice across IP, then all you need to do is to write the application to use the TCP/IP protocols to send data across the network in a way that will make your application run.

This design embeds a value that encourages innovation in applications for the network. It does so both because it minimizes the costs of developing new applications (you don't need the hassle of asking or clearing permission with anyone) and because it avoids strategic behavior by the network owner. Consider again the idea of developing a Voice-over-IP application. If the network is owned by the telephone companies, they would not be excited about



an application that will cannibalize their telephone market. Thus, if permission were required before the VOIP application could be deployed, we might well expect the VOIP application not to be deployed—either because someone developed it, but it was blocked, or because smart developers knew it was a waste of time to develop it, because it would be blocked. As Susan Crawford describes, “The miraculous growth of the Internet has in large part come from the nondiscrimination against higher levels. . . . Innovators at the application layer have been able to assume the continued stable existence of the lower layers.”<sup>55</sup>

The value here is innovation and competition. The network empowers the widest range of innovators—users of the network—and entitles all of them to innovate for this network. Any innovation can be deployed on the network (so long as it respects the TCP/IP protocols). If users of the network like the innovation, then the innovation is a success.

Simultaneously—at least so long as the e2e principle is respected—this design disables the potentially most powerful actor in the network, the network owner, from interfering with the opportunity for innovation within the network. The network owner might not like the stuff being developed, but e2e disables the opportunity to block that development.

In the same way that the original TCP/IP network could be effectively changed so that “gaps” in information about that network could be closed, the TCP/IP network could be changed to remove its e2e character. Indeed, the very tools that I described in Chapter 4 could have this effect. For example, a network owner could scan the packets that were traveling across its network and block any packet that didn’t come from a known, or approved, application. To get on that list, application developers would have to contact the network owner and ask to be included on the list. That change to the way the Internet functions is completely technically possible. Indeed, versions of it are being pursued for both competitive and security reasons. That is, some networks, keen to control the kind of applications that run on the network for competitive reasons, could use this to block disfavored applications (again, think of telephone companies blocking VOIP). Others, keen to avoid viruses or other trouble on their network, could simply decide to block everything to make life simple. Either reason would produce the same result: that innovation on the Internet would be stifled.

As with the stories about “cyberspace,” this case about the Internet also demonstrates the link between architecture and policy. End-to-end is a paradigm for technology that embeds values. Which architecture we encourage is a choice about which policy we encourage. This is true even in the context in which the Internet is not a “place”—even where, that is, it is “just” a medium.



## HOW ARCHITECTURES MATTER AND SPACES DIFFER

The spaces I have described here are different. In some places there is community—a set of norms that are self-enforcing (by members of the community). Features such as visibility (as opposed to anonymity) and nontransience help create those norms; anonymity, transience, and diversity make it harder to create community.

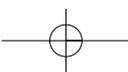
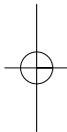
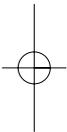
In places where community is not fully self-enforcing, norms are supplemented by rules imposed either through code or by the relevant sovereign. These supplements may further some normative end, but at times they can be in tension with the goal of community building.

If we had to simplify this diversity of spaces by finding a dimension along which we could rank them, one such dimension might be each group's amenability to control. Some groups on this list can be controlled only through norms—.law.cyber, for example. The only technology for changing behavior there—given my commitment not to monitor and punish bad behavior—was the norms of the students in the law school class. Other groups are amenable to other technologies of control. Indeed, as we move from .law.cyber to CC to LambdaMOO to AOL to Second Life, the ability to use these other technologies of control increases, though, of course, that ability is constrained by competition. If the code makes the place no longer attractive, people will leave.

Thus, in CC and AOL, the architects could use technology to change behavior. But if the change is too far removed from what most members think the space is about, members may simply leave. The threat of that constraint turns upon the alternatives, of course. As blogs have flourished, a space like CC would have relatively little market power. AOL's market power is more complicated. There are many alternative ISPs, of course. But once you're a member of one, the costs of migrating are significant.

In LambdaMOO the story is even more complicated. Nothing really binds people to a particular MOO. (There are hundreds, and most are free.) But because characters in a MOO are earned rather than bought, and because this takes time and characters are not fungible, it becomes increasingly hard for members of a successful MOO to move elsewhere. They have the right to exit, but in the sense that Soviet citizens had the right to exit—namely, with none of the assets they had built in their particular world.

Finally, Second Life offers the potential for the most control. Code regulates experience in Second Life more than in any of the other four spaces, and the intimacy of experience in Second Life pulls people into the space and makes escape costly. Again, there are limits to the control, but the controls are



more finely articulated here than in any of the other contexts. And if Philip Rosedale, the CEO of Second Life, is to be believed, the control through code here will only become more subtly expressed. As he described to me:

[O]ur feeling is . . . that we should aggressively move into code anything we can, because of the enhanced scalability it gives us. And we should execute policy outside of code only when absolutely necessary or unfeasible. There are things where we look at them and we say, “Well, we’ll be able to do that in code some day, but for today, we’re just going to do it by hand.”<sup>56</sup>

#### REGULATING CODE TO REGULATE BETTER

I’ve surveyed a range of cyberspaces to make clear the elements of regulation within each. One increasingly important element is code. In cyberspace in particular, but across the Internet in general, code embeds values. It enables, or not, certain control. And as has been the focus of this part, it is also a tool of control—not of government control, at least in the cases I’ve surveyed—but instead control to the end of whatever sovereign does the coding.

These stories suggest a technique, and once we see the idea, we’ll recognize the technique in many different contexts of regulation. If Second Life can use code to better control behavior, what about first-life? If AOL can use code to better control fraud, what about America off-line? If the Internet can use the design of e2e to better enable competition, what does that teach regulators on the ground? How do these techniques of policy inform the practice of policy makers?

The answer is that policy makers have done the same in real space for a long time. Just as Chapter 5 described regulators using code to make behavior more regulable, so too have regulators used code to directly control behavior. Consider a few obvious examples:

#### Tapes

The most significant feature of digital media is that copies can be perfect. Digital media is just data, and data is just a string of 1’s and 0’s. Computers have complex algorithms to verify that when they’ve copied a string of data they’ve copied that string precisely.

This feature thus creates a new risk for sellers of content. While the code of analog copying technology meant that a copy was a degraded version of the original, the code of digital technologies means that a copy could be identical

to the original. That means the threat to content providers from “copies” is greater in the digital world than in the analog world.

Digital Audio Technology (DAT) was the first technology to expose this risk. Like any digital recording, it can, in principle, copy content perfectly. Content providers were thus terrified that piracy from DAT tapes would destroy their industry, so they lobbied Congress effectively to add new laws to protect them from the digital threat.

Congress could have responded to their request in any number of ways. It could have used law to regulate behavior directly, by increasing the penalty for illegal copying. It could have funded a public ad campaign against illegal copying or funded programs in schools to discourage students from buying pirated editions of popular recordings. Congress could have taxed blank tapes and then transferred the revenue to owners of copyrighted material.<sup>57</sup> Or Congress could have tried to regulate DAT technology to weaken the threat that technology presented for copyright.

Congress chose the latter two. The Audio Home Recording Act both taxed blank tapes slightly and regulated the code of digital reproduction technologies directly. The Act requires producers of digital recording devices to install a chip in their systems that implements a code-based system to monitor the copies of any copy made on that machine.<sup>58</sup> The chip would allow a limited number of personal copies, but on copies of copies, the quality of the recording would be degraded. Congress in essence required that the code of digital copying be modified to restore the imperfections that were “natural” in the earlier code.

This again is Congress regulating code as a means of regulating behavior—mandating that multiple copies be imperfect as a way to minimize illegal copying. Like the telephone regulation, this regulation succeeds because there are relatively few manufacturers of DAT technology. Again, given a limited target, the government’s regulation can be effective, and the effect of the government’s regulation is to make more regulable the primary targeted behavior—copyright infringement.

#### Televisions

By the mid-1990s, parents’ concern about the effect that violence on television has on their kids had caught the attention of Congress, and Congress responded through legislation. But given the state of First Amendment law, it would have been difficult for Congress to block violence on television directly. Thus, Congress sought a way to block violence on television indirectly. It sought to require that those broadcasting television content tag their content

with labels that signaled the level of violence in the film, and it mandated that the television industry develop a technology to block content on the basis of those labels.

This was the “V-Chip,” mandated as part of the Telecommunications Act of 1996.<sup>59</sup> The V-chip would facilitate the automatic blocking of television broadcasts, based on criteria of content that have not yet been completely determined. The crudest proposals involve something like the Motion Picture Association’s movie rating system; the more sophisticated envision selections based on a much richer set of factors.

This again is Congress regulating code to affect a targeted behavior (providing violent programming) rather than regulating that behavior directly. The constraint on direct regulation here is similarly a regulability problem. But the lack of regulability in this context comes from constitutional limits, not the inability to track those being regulated by the technology. The constraint of the Constitution thus pushed Congress to require technology to empower parents. By giving parents more power to discriminate, Congress indirectly discourages an ill (exposure to violence) that it is constitutionally unable to regulate directly.<sup>60</sup>

#### Anti-Circumvention

Whatever problem the content industry had with DAT tapes, no doubt they look tiny compared with the problems the content industry has with digital content and the Internet. Although DAT makes perfect copies possible, it doesn’t make distributing those perfect copies any easier. That honor fell to the Internet. Now digital technology not only assured perfect copies of the original, it also made it trivial to distribute those digital copies for free.

As I describe more in Chapter 10, one response to this “feature” of digital technologies is “digital rights management” technology. DRM technologies add code to digital content that disables the simple ability to copy or distribute that content—at least without the technical permission of the DRM technology itself.

Thus, the songs I’ve purchased and downloaded from Apple’s iTunes music store are protected by Apple’s “fairplay” DRM technology. That technology permits me to copy the song to a limited number of machines, but it restricts my ability to copy those songs broadly.

This restriction is effected through code. The “copy” function is produced through code; the DRM technology modifies, or qualifies, that “copy” functionality. It is thus a classic example of code being deployed to restore control over something that (different) code had disabled.

These systems of DRM are privately created. But in 1998, they got an important subsidy of protection from Congress. In the Digital Millennium Copyright Act, Congress banned the creation and distribution of technologies “produced for the purpose of circumventing a technological measure that effectively controls access” to a copyrighted work, or “primarily designed or produced for the purpose of circumventing protection afforded by a technological measure that effectively protects a right of a copyright owner.”<sup>61</sup> By banning this code, Congress aimed to add support to the code content creators were distributing to protect their content. Thus, by directly regulating code, Congress indirectly regulated copyright infringement.

Since this enactment, there has been no end to trouble and litigation surrounding it. Beginning in 1999, the DVD-Copy Control Association began suing individuals and websites that facilitated access to a program, DeCSS, which could be used to decrypt data on DVDs.<sup>62</sup> In July 2001, 27-year-old Russian programmer Dmitry Sklyarov was arrested while giving a presentation in Las Vegas because the company he worked for in Russia had produced software that enabled people to circumvent the access protection technologies built into Adobe’s eBook system.<sup>63</sup> Sklyarov spent six months in an American jail before he was permitted to return to his family in Russia.

The effect of this regulation is hard to measure. The Electronic Frontier Foundation has cataloged its view of the law’s effect five years after the law was enacted.<sup>64</sup> And while the EFF’s view may not be universal, there is a fairly universal surprise at the range of cases that have been brought under the statute. (I doubt the framers of the DMCA imagined that garage door companies would be suing to protect their automatic door openers from competition under the DMCA (they lost).<sup>65</sup>)

#### Broadcast Flags

As broadcast television moves to digital television, copyright holders have become concerned about the risk they face in broadcasting copyrighted content. Unlike an ordinary television broadcast, the quality of a digital broadcast is perfect, so copies of digital broadcasts could likewise be perfect. And the spread of perfect copies of digital broadcasts on a free digital network (the Internet) terrifies copyright holders.

Their response is similar to the response with DAT technologies. First in the FCC, and now in Congress, copyright holders have pushed the government to mandate that any technology capable of reproducing digital broadcasts be architected to respect a “broadcast flag.” If that flag was turned on,

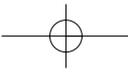
then the technology would be required to block any copy of that content. The content could be played, but it couldn't be reproduced. As Susan Crawford describes it,

The broadcast flag rule, distilled to its essence, is a mandate that all consumer electronics manufacturers and information technology companies ensure that any device that touches digital television content “recognized and give effect to” the flag by protecting content against unauthorized onward distribution. The FCC claimed that the rule would protect digital television (“DTV”) broadcasts from massive redistribution over the Internet.<sup>66</sup>

There is a lot to say about the broadcast flag, and if I were doing the saying, most of it would be bad.<sup>67</sup> But for our purposes, it is the form, not substance, of the broadcast flag that is relevant. This is the most direct example of a regulation of code designed to control primary behavior: law regulating code to make behavior better.

In each case, the government directs an intermediary that has some power over code to change that code to effect a change in behavior. Whether that change in code will effect a change in behavior depends upon the power of the particular application. If the application is a MOO, or an online discussion space like Counsel Connect, the power to control behavior is significantly limited. If the application is AOL or Second Life, the exit costs for a user could well be higher. The scope for effective regulation will thus be greater. And if the application is the Internet, or any digital technology produced or sold in the United States, then the power of the regulator is greater still. Code becomes law even if there remains a capacity to escape the regulation of that code.

These examples point to a general question about how regulation will function. That general point requires many significant qualifications. To understand the effect of code requirements on any regulatory policy will require, as Polk Wagner writes, an understanding that is “profoundly dynamic.”<sup>68</sup> Part of that dynamic, of course, is resistance. Individuals can act to resist the force of code directly. Or individuals can act to resist the force of code through code. As Tim Wu has rightly described, code itself is not necessarily regulation enhancing—code can be used to foil regulation. A gun is a bit of code. It works wonders to destroy the peace. Circumvention technologies are code. They weaken rules reinforcing control. P2P filesharing protocols are code. They undermine the effectiveness of copyright regulations that restrict the freedom to distribute copyrighted works. Whether a particular



regulation will be effective, then, requires consideration of these interactions, and any code-based resistance it might engender. As Wu puts it,

The reason that code matters for law at all is its capability to define behavior on a mass scale. This capability can mean constraints on behavior, in which case code regulates. But it can also mean shaping behavior into legally advantageous forms.<sup>69</sup>

In this second sense, code functions “as an anti-regulatory mechanism: a tool to minimize the costs of law that certain groups will use to their advantage.”<sup>70</sup>

More fundamentally, these complications suggest that a more general framework is needed. I’ve highlighted an interaction between technology, policy, and the law in this chapter. That interaction suggests a much broader model. In the next chapter, I describe that model. In the chapter following that, we will return to the dynamic of code regulation to consider one other important qualification.

