

IO

An Environmentalism for Information

Over the last fifteen years, a group of scholars have finally persuaded economists to believe something noneconomists find obvious: “behavioral economics” shows that people do not act as economic theory predicts. But hold your cheers. This is not a vindication of folk wisdom over the pointy-heads. The deviations from “rational behavior” are not the wonderful cornucopia of human motivations you might imagine. There are patterns. For example, we are systematically likely to overestimate chances of loss and underestimate chances of gain, to rely on simplifying heuristics to frame problems even when those heuristics are contradicted by the facts.

Some of the patterns are endearing; the supposedly “irrational” concerns for distributive equality that persist in all but the economically trained and the extreme right, for example. But most of them simply involve the mapping of cognitive bias. We can take advantage of those biases, as those who sell us ludicrously expensive and irrational warranties on consumer goods do. Or we can correct for them, like a pilot who is trained to rely on his instruments rather than his faulty perceptions when flying in heavy cloud.

This book has introduced you to the wonders and terrors of intellectual property law—the range wars of the Internet age. There have been discussions of synthetic biology and musical sampling, digital locks and the hackers who break them, Jefferson and Macaulay, and the fight over video recorders. Now it is time to sum up.

I would argue that the chapters in this book present evidence of another kind of cognitive bias, one that the behavioral economists have not yet identified. Call it the openness aversion. Cultural agoraphobia. We are systematically likely to undervalue the importance, viability, and productive power of open systems, open networks, and nonproprietary production.

CULTURAL AGORAPHOBIA?

Test yourself on the following questions. In each case, it is 1991 and I have removed from you all knowledge of the years since then. (For some, this might be a relief.)

The first question is a thought experiment I introduced in Chapter 4. You have to design an international computer network. One group of scientists describes a system that is fundamentally open: open protocols and open systems so that anyone could connect to the system and offer information or products to the world. Another group—scholars, businesspeople, bureaucrats—points out the problems. Anyone could connect to the system! They could do anything! The system itself would not limit them to a few approved actions or approved connections. There would be porn, and piracy, and viruses, and spam. Terrorists could put up videos glorifying themselves. Your neighbor's site could compete with the *New York Times* or the U.S. government in documenting the war in Iraq. Better to have a well-managed system in which official approval is required to put up a site, where only a few selected actions are permitted by the network protocols, where most of us are merely recipients of information, where spam, viruses, and piracy (and innovation and participatory culture and anonymous speech) are impossible. Which network design would you have picked? Remember, you have no experience of blogs, or mashups, or Google; no experience of the Web. Just you and your cognitive filters.

Imagine a form of software which anyone could copy and change, created under a license which required subsequent programmers to offer their software on the same terms. Imagine legions of programmers worldwide contributing their creations back into a “commons.” Is this anarchic-sounding method of production economically viable? Could it successfully compete

with the hierarchically organized corporations producing proprietary, closed code, controlled by both law and technology? Be truthful.

Finally, set yourself the task of producing the greatest reference work the world has ever seen. You are told that it must cover everything from the best Thai food in Durham to the annual rice production of Thailand, from the best places to see blue whales to the history of the Blue Dog Coalition. Would you create a massive organization of paid experts, each assigned a topic, with hierarchical layers of editors above them, producing a set of encyclopedic tomes that are rigorously controlled by copyright and trademark? Or would you wait for hobbyists, governments, scientists, and volunteer encyclopedists to produce, and search engines to organize and rank, a cornucopia of information? I know which way I would have bet in 1991. But I also know that the last time I consulted an encyclopedia was in 1998. You?

It is not that openness is always right. It is not. Often we need strong intellectual property rights, privacy controls, and networks that demand authentication. Rather, it is that we need a balance between open and closed, owned and free, and we are systematically likely to get the balance wrong. (How did you do on the test?) Partly this is because we still don't understand the kind of property that lives on networks; most of our experience is with tangible property. Sandwiches that one hundred people cannot share. Fields that can be overgrazed if outsiders cannot be excluded. For that kind of property, control makes more sense. Like astronauts brought up in gravity, our reflexes are poorly suited for free fall. Jefferson's words were true even of grain elevators and hopper-boys. But in our world, the proportion of intangible to tangible property is much, much higher. The tendency to conflate intellectual and real property is even more dangerous in a networked world. We need his words more than he did.

Each of the questions I asked is related to the World Wide Web. Not the Internet, the collective name for the whole phenomenon, including the underlying methods of sending and receiving packets. Some version of the underlying network has been around for much longer, in one form or another. But it only attracted popular attention, only revolutionized the world, when on top of it was built the World Wide Web—the network of protocols and pages and hyperlinks that is so much a part of our lives and which arose only from Tim Berners-Lee's work at CERN in 1991.

My daughter will graduate from college in the year 2011. (At least, we both hope so.) She is older than the Web. It will not even have had its twentieth birthday on her graduation day. By Christmas of 2012, it will be able to drink

legally in the United States. I wrote those sentences, but I find it hard to believe them myself. A life without the Web is easy to remember and yet hard to recapture fully. It seems like such a natural part of our world, too fixed to have been such a recent arrival, as if someone suggested that all the roads and buildings around you had arrived in the last fifteen years.

Some of you may find these words inexplicable because you live in a happy, Thoreau-like bliss, free of any contact with computer networks. If so, I take my hat off to you. The world of open sky and virtuous sweat, of books and sport and laughter, is no less dear to me than to you. Having an avatar in a virtual world holds the same interest as elective dental surgery. I care about the Web not because I want to live my life there, but because of what it has allowed us to achieve, what it represents for the potential of open science and culture. That, I think, is something that Thoreau (and even Emerson for that matter) might have cared about deeply. Yet, as I suggested earlier in this book, I seriously doubt that we would create the Web today—at least if policy makers and market incumbents understood what the technology might become early enough to stop it.

I am not postulating some sinister “Breakages, Limited” that stifles technological innovation. I am merely pointing out the imbalance between our intuitive perceptions of the virtues and dangers of open and closed systems, an imbalance I share, quite frankly.

In place of what we have today, I think we would try, indeed we are trying, to reinvent a tamer, more controlled Web and to change the nature of the underlying network on which it operates. (This is a fear I share with those who have written about it more eloquently than I, particularly Larry Lessig and Yochai Benkler.) We would restrict openness of access, decrease anonymity, and limit the number of actions that a network participant could perform. The benefits would be undeniable. It would cut down on spam, viruses, and illicit peer-to-peer file sharing. At the same time, it would undercut the iconoclastic technological, cultural, and political potential that the Web offers, the ability of a new technology, a new service to build on open networks and open protocols, without needing approval from regulators or entrenched market players, or even the owners of the Web pages to which you link.

Imagine, by contrast, an Internet and a World Wide Web that looked like America Online, circa 1996, or CompuServe, or the French state network Minitel. True, your exposure to penis-enhancement techniques, misspelled stock tips, and the penniless sons of Nigerian oil ministers would be reduced. That sounds pretty attractive. But the idea that the AOL search engine would

be replaced by Yahoo and then Google, let alone Google Maps? That new forms of instant messaging would displace CompuServe's e-mail? That the Chinese dissident would have access to anonymized Internet services, that you might make phone calls worldwide for free from your computer, or that a blog like BoingBoing would end up having more page views than many major newspapers? Forget it. Goodbye to the radical idea that anyone can link to any page on the network without permission. A revised network could have the opposite rule and even impose it by default.

A tamer network could keep much tighter control over content, particularly copyrighted content. You might still get the video of the gentlemen doing strange things with Mentos and soda bottles, though not its viral method of distribution. But forget about "George Bush Doesn't Care About Black People" and all your favorite mashups. Its controlled network of links and its limited access would never unleash the collective fact-gathering genius the Web has shown. For a fee, you would have Microsoft Encarta and the *Encyclopedia Britannica* online. What about the "right-click universe" of knowledge about the world gathered by strangers, shared on comparatively open sites worldwide, and ordered by search engines? What about Wikipedia? I think not.

The counterfactual I offer is not merely a counterfactual. Yes, we got the Web. It spread too fast to think of taming it into the more mature, sedate "National Information Infrastructure" that the Clinton administration imagined. But as Larry Lessig pointed out years ago, the nature of a network can always be changed. The war over the control and design of the network, and the networked computer, is never-ending. As I write these words, the battles are over "trusted computing" and "Net neutrality." Trusted computing is a feature built into the operating system which makes it impossible to run processes that have not been approved by some outside body and digitally identified. It would indeed help to safeguard your computer from viruses and other threats and make it harder to copy material the content owners did not want you to copy (perhaps even if you had a right to). In the process it would help to lock in the power of those who had a dominant position in operating systems and popular programs. (Microsoft is a big supporter.) It would make open source software, which allows users to modify programs, inherently suspect. It would, in fact, as Jonathan Zittrain points out, change the nature of the general-purpose computer, which you can program to do *anything*, back toward the terminal which tells *you* what functions are allowed.¹ Think of a DVD player.

The attack on Net neutrality, by contrast, is an attempt by the companies who own the networks to be allowed to discriminate between favored and disfavored content, giving the former preferential access. (One wit analogized it to letting the phone company say, “we will delay your call to Pizza Hut for sixty seconds, but if you want to be put through to our featured pizza provider immediately, hit nine now!”) Taken together, these proposals would put the control of the computer back in the hands of the owners of the content and the operating system, and control of the network users’ choices in the hands of the person who sells them their bandwidth. At the same time, our intellectual property agenda is filled with proposals to create new intellectual property rights or extend old ones. That is the openness aversion in action.

Now, perhaps to you, the closed alternatives still sound better. Perhaps you do not care as much about the kind of technological dynamism, or anonymous speech, or cultural ferment that thrills the digerati. Perhaps you care more about the risks posed by the underlying freedom. That is a perfectly reasonable point of view. After all, openness does present real dangers; the same freedom given to the innovator, the artist, and the dissident is given to the predator and the criminal. At each moment in history when we have opened a communications network, or the franchise, or literacy, reasonable people have worried about the consequences that might ensue. Would expanded literacy lead to a general coarsening of the literary imagination? (Sometimes, perhaps. But it would and did lead to much more besides, to literature and culture of which we could not have dreamed.) Would an expanded franchise put the control of the state into the hands of the uneducated? (Yes, unless we had free national educational systems. “Now we must educate our masters” was the slogan of the educational reformers after the enlargement of the franchise in Britain in the nineteenth century. Openness sometimes begets openness.) Would translating the Bible from Latin into the vernacular open the door to unorthodox and heretical interpretations, to a congregation straying because they did not need to depend on a priestly intermediary with privileged access to the text? (Oh, yes indeed.) Would TV and radio play into the hands of demagogues? (Yes, and help expose their misdeeds.)

Openness is not always right. Far from it. But our prior experience seems to be that we are systematically better at seeing its dangers than its benefits. This book has been an attempt, in the sphere of intellectual property, to help us counteract that bias. Like the pilot in the cloud looking at his instruments, we might learn that we are upside down. But what do we do about it?

LEARNING FROM ENVIRONMENTALISM

I have argued that our policies are distorted not merely by industry capture or the power of incumbent firms, but by a series of cultural and economic biases or presuppositions: the equation of intellectual property to physical property; the assumption that whenever value is created, an intellectual property right should follow; the romantic idea of creativity that needs no raw material from which to build; the habit of considering the threats, but not the benefits, of new technologies; the notion that more rights will automatically bring more innovation; the failure to realize that the public domain is a vital contributor to innovation and culture; and a tendency to see the dangers of openness, but not its potential benefits.²

One of the most stunning pieces of evidence to our aversion to openness is that, for the last fifty years, whenever there has been a change in the law, it has almost always been to expand intellectual property rights. (Remember, this implies that every significant change in technology, society, or economy required more rights, never less, nor even the same amount.) We have done all this almost entirely in the absence of empirical evidence, and without empirical reconsideration to see if our policies were working. As I pointed out in the last chapter, intellectual property policy is an “evidence-free zone.” It runs on faith alone and its faith consists of the cluster of ideas I have outlined in this book. Whether we call this cluster of ideas maximalism, cultural agoraphobia, or the openness aversion, it exercises a profound influence on our intellectual property and communications policy.

These ideas are not free-floating. They exist within, are influenced by, and in turn influence, a political economy. The political economy matters and it will shape any viable response. Even if the costs of getting the policies wrong are huge and unnecessary—think of the costs of the copyright extensions that lock up most of twentieth-century culture in order to protect the tiny fraction of it that is still commercially available—they are spread out over the entire population, while the benefits accrue to a small group of commercial entities that deeply and sincerely believe in the maximalist creed. This pattern of diffuse but large losses and concentrated gains is, as Mancur Olson taught us, a recipe for political malfunction.³ Yet the problem is even deeper than that—in four ways.

First, though intellectual property rules will profoundly shape science, culture, and the market in the information age, they just seem obscure, wonkish, hard to get excited about. Certainly, people can get upset about individual

examples—overbroad patents on human genes, copyright lawsuits against whistleblowers who leak e-mails showing corporate misdeeds that threaten the integrity of electronic voting, rules that paralyze documentary filmmakers, or require payment for sampling three notes from a prior song, extensions of rights that allow patents on auctions or business methods, make genres such as jazz seem legally problematic, create new rights over facts, or snarl up foundational technologies. But they see each of these as an isolated malfunction, not part of a larger social problem or set of attitudes.

Second, what holds true for issues, also holds true for communities. What links the person writing open source software, and trying to negotiate a sea of software patents in the process, to the film archivist trying to stir up interest in all the wonderful “orphan films”—still under copyright but with no copyright owner we can find—before they molder away into nitrate dust? When a university collaborates with Google to digitize books in their collection for the purposes of search and retrieval, even if only a tiny portion of the text will be visible for any work still under copyright, does it sense any common interest with the synthetic biologist trying to create the BioBricks Foundation, to keep open the foundational elements of a new scientific field? Both may be sued for their efforts—one connection at least.

When a developing nation tries to make use of the explicit “flexibilities” built into international trade agreements so as to make available a life-saving drug to its population through a process of compulsory licensing and compensation, it will find itself pilloried as a lawbreaker—though it is not—or punished through bilateral agreements. Will that process form any common interest with the high-technology industries in the United States who chafe at the way that current intellectual property rules enshrine older technologies and business methods and give them the protection of law? There are some links between those two situations. Will the parties see those links, or will the developing world’s negotiators think that the current intellectual property rules express some monolithic “Western” set of interests? Will the high-tech companies think this is just an issue of dumb lawyers failing to understand technology? Each gap in understanding of common interest is a strike against an effective response.

Third, an effective political response would actually be easier if our current rules came merely from the relentless pursuit of corporate self-interest. (Here I part company with those who believe that self-interest is simply “there”—not shaped by socially constructed ideas, attitudes, ideologies, or biases.) In fact, the openness aversion sometimes obscures self-interest as well as the

public interest. Think of the relentless insistence of the movie companies on making video recorders illegal. Nor does the framework of maximalism help if our goal is to have all the interested economic actors in the room when policy is made. For example, by framing issues of communications policy or Internet regulation as questions of intellectual property, we automatically privilege one set of interested parties—content owners—over others who also have a large economic stake in the matter.

Fourth, and finally, the biggest problem is that even if one could overcome the problems of political interest, or ideological closed-mindedness, the answers to many of these questions require balance, thought, and empirical evidence—all qualities markedly missing in the debate. If the answer were that intellectual property rights are bad, then forming good policy would be easy. But that is as silly and one-sided an idea as the maximalist one I have been criticizing here. Here are three examples:

1. Drug patents do help produce drugs. Jettisoning them is a bad idea—though experimenting with additional and alternative methods of encouraging medical innovation is a very good one.
2. I believe copyrights over literary works should be shorter, and that one should have to renew them after twenty-eight years—something that about 85 percent of authors and publishers will not do, if prior history is anything to go by. I think that would give ample incentives to write and distribute books, and give us a richer, more accessible culture and educational system to boot, a Library of Congress where you truly can “click to get the book” as my son asked me to do years ago now. But that does not mean that I wish to abolish copyright. On the contrary, I think it is an excellent system.
3. All the empirical evidence shows that protecting compilations of facts, as the European Database Directive does, has been a profound failure as a policy, imposing costs on consumers without encouraging new database production. But if the evidence said the opposite, I would support a new database right.

We need a political debate about intellectual property that recognizes these trade-offs; that does not impose simplistic, one-sided solutions; that looks to evidence. We need to understand the delicate and subtle balance between property and the opposite of property, the role of rights, but also of the public domain and the commons. Building a theory, let alone a movement, around such an issue is hard. Doing so when we lack some of the basic

theoretical tools and vocabularies is daunting. We do not even have a robust conception of the public domain. If they think of it as a legal issue at all, people simply think of it as whatever is left over after an endless series of rights have been carved out. Can one build a politics to protect a *residue*?

So we have at least four problems: an issue that is perceived as obscure, affecting scattered groups with little knowledge of each other's interest, dominated by an ideology that is genuinely believed by its adherents, in the place of which we have to make careful, balanced, empirically grounded suggestions. Assume for a moment the need for a politics of intellectual property that seeks a solution to these four problems. What might such a politics look like?

I have argued that in a number of respects, the politics of intellectual property and the public domain is at the stage that the American environmental movement was at in the 1950s. In 1950, there were people who cared strongly about issues we would now identify as "environmental"—supporters of the park system and birdwatchers, but also hunters and those who disdained chemical pesticides in growing their foods. In the world of intellectual property, we have start-up software engineers, libraries, appropriationist artists, parodists, biographers, and biotech researchers. In the 50s and 60s, we had flurries of outrage over particular crises—burning rivers, oil spills, dreadful smog. In the world of intellectual property, we have the kind of stories I have tried to tell here. Lacking, however, is a general framework, a perception of common interest in apparently disparate situations.

Crudely speaking, the environmental movement was deeply influenced by two basic analytical frameworks. The first was the idea of ecology: the fragile, complex, and unpredictable interconnections between living systems. The second was the idea of welfare economics—the ways in which markets can fail to make activities internalize their full costs.⁴ The combination of the two ideas yielded a powerful and disturbing conclusion. Markets would *routinely* fail to make activities internalize their own costs, particularly their own environmental costs. This failure would, routinely, disrupt or destroy fragile ecological systems, with unpredictable, ugly, dangerous, and possibly irreparable consequences. These two types of analysis pointed to a general interest in environmental protection and thus helped to build a large constituency which supported governmental efforts to that end. The duck hunter's preservation of wetlands as a species habitat turns out to have wider functions in the prevention of erosion and the maintenance of water quality. The decision to burn coal rather than natural gas for power generation may have impacts on everything from forests to fisheries. The attempt to reduce greenhouse gases and

mitigate the damage from global warming cuts across every aspect of the economy.

Of course, it would be silly to think that environmental policy was fueled only by ideas rather than more immediate desires. As William Ruckelshaus put it, “With air pollution there was, for example, a desire of the people living in Denver to see the mountains again. Similarly, the people living in Los Angeles had a desire to see one another.” Funnily enough, as with intellectual property, changes in communications technology also played a role. “In our living rooms in the middle sixties, black and white television went out and color television came in. We have only begun to understand some of the impacts of television on our lives, but certainly for the environmental movement it was a bonanza. A yellow outfall flowing into a blue river does not have anywhere near the impact on black and white television that it has on color television; neither does brown smog against a blue sky.”⁵ More importantly perhaps, the technologically fueled deluge of information, whether from weather satellites or computer models running on supercomputers, provided some of the evidence that—eventually—started to build a consensus around the seriousness of global warming.

Despite the importance of these other factors, the ideas I mentioned—ecology and welfare economics—were extremely important for the environmental movement. They helped to provide its agenda, its rhetoric, and the perception of common interest underneath its coalition politics. Even more interestingly, for my purposes, those ideas—which began as inaccessible scientific or economic concepts, far from popular discourse—were brought into the mainstream of American politics. This did not happen easily or automatically. Popularizing complicated ideas is hard work. There were popular books, television discussions, documentaries on Love Canal or the California kelp beds, op-ed pieces in newspapers, and pontificating experts on TV. Environmental groups both shocking and staid played their part, through the dramatic theater of a Greenpeace protest or the tweedy respectability of the Audubon Society. Where once the idea of “the Environment” (as opposed to “my lake,” say) was seen as a mere abstraction, something that couldn’t stand against the concrete benefits brought by a particular piece of development, it came to be an abstraction with both the force of law and of popular interest behind it.

To me, this suggests a strategy for the future of the politics of intellectual property, a way to save our eroding public domain. In both areas, we seem to have the same recipe for failure in the structure of the decision-making process. Democratic decisions are made badly when they are primarily made by

and for the benefit of a few stakeholders, whether industrialists or content providers. This effect is only intensified when the transaction costs of identifying and resisting the change are high. Think of the costs and benefits of acid rain-producing power generation or—less serious, but surely similar in form—the costs and benefits of retrospectively increasing copyright term limits on works for which the copyright had already expired, pulling them back out of the public domain. There are obvious benefits to the heirs and assigns of authors whose copyright has expired in having Congress put the fence back up around this portion of the intellectual commons. There are clearly some costs—for example, to education and public debate—in not having multiple, competing low-cost editions of these works. But these costs are individually small and have few obvious stakeholders to represent them.

Yet, as I have tried to argue here, beyond the failures in the decision-making process, lie failures in the way we think about the issues. The environmental movement gained much of its persuasive power by pointing out that for structural reasons we were likely to make bad environmental decisions: a legal system based on a particular notion of what “private property” entailed and an engineering or scientific system that treated the world as a simple, linearly related set of causes and effects. In both of these conceptual systems, the environment actually disappeared; there was no place for it in the analysis. Small surprise, then, that we did not preserve it very well. I have argued that the same is true about the public domain. The confusions against which the Jefferson Warning cautions, the source-blindness of a model of property rights centered on an “original author,” and the political blindness to the importance of the public domain as a whole (not “my lake,” but “the Environment”), all come together to make the public domain disappear, first in concept and then, increasingly, as a reality. To end this process we need a cultural environmentalism, an environmentalism of the mind, and over the last ten years we have actually begun to build one.

Cultural environmentalism is an idea, an intellectual and practical movement, that is intended to be a solution to a set of political and theoretical problems—an imbalance in the way we make intellectual property policy, a legal regime that has adapted poorly to the transformation that technology has produced in the scope of law, and, perhaps most importantly, a set of mental models, economic nostrums, and property theories that each have a public domain-shaped hole at their center.

The comparison I drew between the history of environmentalism and the state of intellectual property policy had a number of facets. The environmental

movement had “invented” the concept of the environment and used it to tie together a set of phenomena that would otherwise seem very separate. In doing so, it changed perceptions of self-interest and helped to form coalitions where none had existed before—just as earth science built upon research into the fragile interconnections of ecology and on the Pigouvian analysis of economic externalities. I argue that we need to make visible the invisible contributions of the public domain, the “ecosystem services” performed by the underappreciated but nevertheless vital reservoir of freedom in culture and science.⁶ And, just as with environmentalism, we need not only a semantic reorganization, or a set of conceptual and analytic tools, but a movement of people devoted to bringing a goal to the attention of their fellow citizens.

I have tried hard to show that there is something larger going on under the realpolitik of land grabs by Disney and campaign contributions by the Recording Industry Association of America. But it would be an equal and opposite mistake to think that this is just about a dysfunctional discourse of intellectual property. In this part of the analysis, too, the environmental movement offers some useful practical reminders. The ideas of ecology and environmental welfare economics were important, but one cannot merely write *A Sand County Almanac* and hope the world will change. Environmentalists piggybacked on existing sources of conservationist sentiment—love of nature, the national parks movement, hikers, campers, birdwatchers. They built coalitions between those who might be affected by environmental changes. They even stretched their political base by discovering, albeit too slowly, the realities of environmental racism, on the one hand, and the benefits of market solutions to some environmental problems on the other. Some of these aspects, at least, could be replicated in the politics of intellectual property.

Ten years ago, when I first offered the environmental analogy, I claimed that intellectual property policy was seen as a contract struck between industry groups—something technical, esoteric, and largely irrelevant to individual citizens, except in that they were purchasers of the products that flowed out of the system. Whether or not that view has ever been tenable, it is not so in a digital age. Instead, I offered the basic argument laid out here—that we needed a “politics of intellectual property” modeled on the environmental movement to create a genuine and informed political debate on intellectual property policy.⁷

So far, I have concentrated on the theoretical and academic tools such a debate would need—focusing particularly on property theory and on economic

analysis and its limits. But if there is to be a genuinely democratic politics of intellectual property, we would need an institutional diversity in the policy-making debate that was comparable to that of the environmental movement.

Environmentalism presents us with a remarkable diversity of organizational forms and missions. We have Greenpeace, the Environmental Legal Defense Fund, groups of concerned scientists, and the Audubon Society, each with its own methods, groups of supporters, and sets of issues. Yet we also have local and pragmatic coalitions to save a particular bit of green space, using the private tools of covenants and contracts.⁸ I think we can see the beginnings of the replication of that institutional diversity in the world of intangible property.

Ten years ago, civil society had little to offer in terms of groups that represented anything other than an industry position on intellectual property, still less ones that took seriously the preservation of the public domain or the idea that intellectual property policy was a matter of balance, rather than simple maximization of rights. There were the librarians and a few academics. That was about it. This position has changed radically.

There are academic centers that concentrate on the theoretical issues discussed in this book—one of them at my university. Thanks in large part to the leadership of Pamela Samuelson, there are law student clinics that do impact litigation on issues such as fair use and that represent underserved clients such as documentarians. But beyond academic work, there are organizations that have dedicated themselves to advocacy and to litigation around the themes of preservation of the public domain, defense of limitations and exceptions in copyright, and the protection of free speech from the effects of intellectual property regulation of both content and the communications infrastructure. The Electronic Frontier Foundation did exist ten years ago, but its coverage of intellectual property issues was only episodic. Its portfolio of litigation and public education on the subject is now nothing short of remarkable. Public Knowledge's valuable lobbying and education is another obvious example. International organizations with similar aims include the Open Rights Group in the United Kingdom.⁹

Organizing has also taken place around particular cases—such as *Eldred v. Ashcroft*, the challenge to the Sonny Bono Copyright Term Extension Act.¹⁰ Activity is not confined to the world of copyright. The Public Patent Foundation combats “patent creep” by exposing and challenging bad patents.¹¹

It would be remiss not to mention the international Access to Knowledge, or A2K, movement, inspired by the work of Jamie Love.¹² While its focus is

on the kinds of issues represented by the access-to-medicines movement, it has made the idea of balance in intellectual property and the protection of the public domain one of its central components. Mr. Love himself is also the central figure behind the idea of a Research and Development Treaty which would amend international trade agreements to make intellectual property merely one of a whole range of economic methods for stimulating innovation.¹³ His work has touched almost every single one of the movements discussed here.

The Access to Knowledge movement has many institutional variants. The Development Agenda at the World Intellectual Property Organization (WIPO), put forward by India and Brazil, includes similar themes, as do the Geneva Declaration and the Adelphi Charter produced by the United Kingdom's Royal Society for the Encouragement of Arts, Manufactures and Commerce.¹⁴ History is full of wordy charters and declarations, of course. By themselves they mean little. Yet the level of public and media attention paid to them indicates that intellectual property policy is now of interest beyond a narrow group of affected industries. To underscore this point, several major foundations have introduced intellectual property initiatives, something that would have been inconceivable ten years ago.¹⁵

Finally, to complete the analogy to the land trust, we have the organizations I mentioned earlier, such as Creative Commons and the Free Software Foundation.¹⁶ The latter group pioneered within software the attempt to create a licensed "commons" in which freedoms are guaranteed. The licensed commons replaces the law's default rules with choices made by individuals, the effects of which are magnified by collective action. The end result is a zone of public freedom enabled by private choice.

If one looks at these institutions and actors and at the range of issues on which they focus—from software to drug patents, from reverse engineering to access to archival records—the obvious question is, how did they overcome the collective action problem? What ties together a critique of digital locks and the access-to-medicines movement? Again, I think the answer points to the usefulness of the environmental analogy. As I pointed out, the invention of the "environment" trope tied together groups whose interests, considered at a lower level of abstraction, seemed entirely different—hunters and bird-watchers, antipollution protesters and conservation biologists. The idea of the "environment" literally created the self-interest or set of preferences that ties the movement together. The same is true here. Apparently disparate interests are linked by ideas of the protection of the public domain and of the

importance of a balance between protection and freedom in cultural and scientific ecology.¹⁷

But even a broad range of initiatives and institutions would not, in and of themselves, produce results. One must convince people that one's arguments are good, one's institutional innovations necessary, one's horror stories disturbing. Environmentalism has managed to win the battle for clarity—to make its points clearly enough that they ceased to be dismissed as “arcane” or technical, to overcome neglect by the media, to articulate a set of concerns that are those of any educated citizen. The other striking phenomenon of the last ten years is the migration of intellectual property issues off the law reviews or business pages and onto the front pages and the editorial pages. Blogs have been particularly influential. Widely read sites such as Slashdot and Boing-Boing have multiple postings on intellectual property issues each day; some are rants, but others are at a level of sophistication that once would have been confined to academic discussion.¹⁸ Scientists passionately debate the importance of open access to scholarly journals. Geographers and climatologists fume over access to geospatial data. The movement has been pronounced enough to generate its own reaction. The popular comics site “xkcd” has strips critical of the Digital Millennium Copyright Act,¹⁹ but also a nerdily idyllic picture of a stick figure reclining under a tree and saying, “Sometimes I just can't get outraged over copyright law.”²⁰ That cartoon now resides on my computer desktop. (It is under a Creative Commons license, ironically enough.)

Who can blame the stick figure? Certainly not I. Is it not silly to equate the protection of the environment with the protection of the public domain? After all, one is the struggle to save a planetary ecology and the other is just some silly argument about legal rules and culture and science. I would be the first to yield primacy to the environmental challenges we are facing. Mass extinction events are to be avoided, particularly if they involve you personally. Yet my willingness to minimize the importance of the rules that determine who owns science and culture goes only so far.

A better intellectual property system will not save the planet. On the other hand, one of the most promising sets of tools for building biofuels comes from synthetic biology. Ask some of the leading scientists in that field why they devoted their precious time to trying to work out a system that would offer the valuable incentives that patents provide while leaving a commons of “biobricks” open to all for future development. I worry about these rules naturally; they were forced to do so. A better intellectual property system certainly

will not end world hunger. Still it is interesting to read about the lengthy struggles to clear the multiple, overlapping patents on *GoldenRice*TM—a rice grain genetically engineered to cure vitamin deficiencies that nearly perished in a thicket of blurrily overlapping rights.²¹

A better intellectual property system will not cure AIDS or rheumatoid arthritis or Huntington's disease or malaria. Certainly not by itself. Patents have already played a positive role in contributing to treatments for the first two, though they are unlikely to help much on the latter two; the affected populations are too few or too poor. But overly broad, or vague, or confusing patents could (and I believe have) hurt all of those efforts—even those being pursued out of altruism. Those problems could be mitigated. Reforms that made possible legal and facilitated distribution of patented medicines in Africa might save millions of lives. They would cost drug companies little. Africa makes up 1.6 percent of their global market. Interesting alternative methods have even been suggested for encouraging investment in treatments for neglected diseases and diseases of the world's poor. At the moment, we spend 90 percent of our research dollars on diseases that affect 10 percent of the global population. Perhaps this is the best we can do, but would it not be nice to have a vigorous public debate on the subject? Some possible innovations are much easier. A simple rule that required the eventual free publication online of all government-funded health research, under open licenses, rather than its sequestration behind the paywalls of commercial journals, could help fuel remarkable innovations in scientific synthesis and computer-aided research while giving citizens access to the research for which they have already paid.

Good intellectual property policy will not save our culture. But bad policy may lock up our cultural heritage unnecessarily, leave it to molder in libraries, forbid citizens to digitize it, even though the vast majority of it will never be available publicly and no copyright owner can be found. Would you not prefer the world in which your children could look at the Library of Congress online catalogue and click to get the book or film or song that otherwise languished as an "orphan work"? Good intellectual policy will not necessarily give us great new music. But the policy we have today would make some of the music we most cherish illegal, or at least legally questionable. Does that inspire confidence for the future? As for the World Wide Web, I offer again my thought experiment from the first part of this chapter. Would we be more likely to invent it or forbid it today? We are certainly working busily to change the openness of the general-purpose computer, the neutrality of the network, and the degree of control that content companies can exert over hardware.

I do not claim that the issues I have written about here are the most important problem the world faces. That would be ridiculous. But I do claim that they are facets of a very important problem and one to which we are paying far too little attention.

I would also be the first to admit that these issues are complicated. Even if we heeded the precepts I have outlined in this book, even if we actually started to look at intellectual property as an empirical question, even if we turned to data rather than faith for our assessments, reasonable people would disagree about much. Some of the most ludicrous recent excesses—huge retrospective copyright term extensions, database rights, proposed webcasting treaties, business method patents—do not pass the laugh test, in my view and that of most scholars. Stopping and then reversing that tide would be valuable, even transformative, but other issues are a closer call.

It is also true that we do not have all the tools we need. A lot remains to be done, both academically and practically. We need better evidence. We need property theories that give us as rich a conception of property's outside—of the public domain and the commons—as we have of property itself. We need to rethink some of our policies of international harmonization and reconsider what types of policy actually benefit the developing world. We should explore ways of compensating artists that are very different from the ones we use now, and study the use of distributed creativity and open source in new areas of science and culture.

Difficulties aside, I have tried here to show that we need a cultural environmental movement, a politics that enables us first to see and then to preserve the public domain, to understand its contributions to our art, our technology, and our culture. Where is that movement now?

There is cause for both concern and optimism. Concern, because it is still hard for courts, legislators, policy makers, and citizens to see beyond the word “property” to the reality underneath. I started this book with the question from my son about the online catalogue of the Library of Congress: “Where do you click to get the book?” In 2003 the Supreme Court heard *Eldred v. Ashcroft*, the challenge to retrospective copyright term extension. Over two strong dissents, the Court upheld the constitutionality of the act against both First Amendment and Copyright Clause challenges. The dead had their copyrights extended yet again. The widest legal restriction of speech in the history of the Republic—putting off-limits most twentieth-century books, poems, films, and songs for another twenty years without a corresponding speech benefit or incentive—can proceed without significant First Amendment review.

Does such a decision mean the task this book undertakes—to take seriously the contributions of the public domain to innovation, culture, and speech—is ultimately doomed, whatever its intellectual merits, to face a hostile or uncomprehending audience? Admittedly, *Eldred* focused specifically on two particular constitutional claims. Still, the attitude of the majority toward the importance of the public domain—whether in the textual limitations on Congress’s power or the application of the First Amendment—can hardly be cause for optimism. And yet . . . The media reaction was remarkable.

The *New York Times* was sufficiently unfamiliar with the term “public domain” that it was not entirely sure whether or not to use the definite article in front of it. But unfamiliarity did not imply complacency. An editorial declared that this decision “makes it likely that we are seeing the beginning of the end of public domain and the birth of copyright perpetuity. Public domain has been a grand experiment, one that should not be allowed to die. The ability to draw freely on the entire creative output of humanity is one of the reasons we live in a time of such fruitful creative ferment.”²² The *Washington Post*, though more inclined to agree that retrospective extension might be constitutional, declared the copyright system to be “broken” in that it “effectively and perpetually protects nearly all material that anyone would want to cite or use. That’s not what the framers envisioned, and it’s not in the public interest.”²³

I could not agree more. But as I have tried to show here, the process is not limited to copyright, or culture, or texts, or the United States. Think of the stories about business method patents, or synthetic biology, or the regulation of musical borrowing on the atomic level. Think of the discussion of the openness aversion that began this chapter. In the middle of the most successful and exciting experiment in nonproprietary, distributed creativity in the history of the species, our policy makers can see only the threat from “piracy.” They act accordingly. Our second enclosure movement is well under way. The poem with which I began Chapter 3 told us: “And geese will still a common lack / Till they go and steal it back.” I cannot match the terseness or the rhyme, but if we assume that the enclosure of the commons of the mind will bring us prosperity, great science, and vibrant culture, well, we will look like very silly geese indeed.