

8

A Creative Commons

If you go to the familiar Google search page and click the intimidating link marked “advanced search,” you come to a page that gives you more fine-grained control over the framing of your query. Nestled among the choices that allow you to pick your desired language, or exclude raunchy content, is an option that says “usage rights.” Click “free to use or share” and then search for “physics textbook” and you can download a 1,200-page physics textbook, copy it, or even print it out and hand it to your students. Search for “Down and Out in the Magic Kingdom” and you will find Cory Doctorow’s fabulous science fiction novel, online, in full, for free. His other novels are there too—with the willing connivance of his commercial publisher. Search for “David Byrne, My Fair Lady” and you will be able to download Byrne’s song and make copies for your friends. You’ll find songs from Gilberto Gil and the Beastie Boys on the same page. No need to pay iTunes or worry about breaking the law.

Go to the “advanced” page on Flickr, the popular photo sharing site, and you will find a similar choice marked “Creative Commons License.” Check that box and then search for “Duke Chapel” and

you will get a selection of beautiful photos of the lovely piece of faux Gothic architecture that sits about three hundred yards from the office where I am writing these words. You can copy those photos, and 66 million others on different subjects, share them with your friends, print them for your wall, and, in some cases, even use them commercially. The same basic tools can be found on a range of specialized search engines with names like OWL Music Search, BlipTV, SpinExpress, and OERCommons. Searching those sites, or just sticking with the advanced options on Google or Yahoo, will get you courses in music theory, moral philosophy, and C++ programming from famous universities; a full-length movie called *Teach* by Oscar-winning director Davis Guggenheim; and free architectural drawings that can be used to build low-cost housing. At the Wellcome Library, you will find two thousand years of medical images that can be shared freely. Searching for “skeleton” is particularly fun. You can even go to your favorite search engine, type in the title of this book, find a site that will allow you to download it, and send the PDF to a hundred friends, warmly anticipating their rapturous enjoyment. (Better ask them first.)

All this copying and sharing and printing sounds illegal, but it is not (at least if you went through the steps I described). And the things you can do with this content do not stop with simply reproducing it, printing it on paper, or sending it by e-mail. Much of it can be changed, customized, remixed—you could rewrite the module of the class and insert your own illustrations, animate the graphs showing calculus in action, morph the photo into something new. If you search for a musician with the unpromising name “Brad Sucks,” you will find a Web site bearing the modest subtitle “A one man band with no fans.” Brad, it turns out, does not suck and has many fans. What makes him particularly interesting is that he allows those fans, or anyone else for that matter, to remix his music and post their creations online. I am particularly fond of the Matterovermind remix of “Making Me Nervous,” but it may not be to your taste. Go to a site called ccMixer and you will find that musicians, famous and obscure, are inviting you to sample and remix their music. Or search Google for Colin Mutchler and listen to a haunting song called “My Life Changed.” Mr. Mutchler and a violinist called Cora Beth Bridges whom he had never met created that song together. He posted a song called “My Life” online, giving anyone the freedom to add to it, and she did—“My Life.” Changed.

On December 15, 2002, in San Francisco, a charitable organization called Creative Commons was launched. (Full disclosure: I have been a proud board

member of Creative Commons since its creation.) Creative Commons was the brainchild of Larry Lessig, Hal Abelson, and Eric Eldred. All the works I have just described—and this book itself—are under Creative Commons licenses. The authors and creators of those works have chosen to share it with the world, with you, under generous terms, while reserving certain rights for themselves. They may have allowed you to copy it, but not to alter it—to make derivative works. Or they may have allowed you to use it as you wish, so long as you do so noncommercially. Or they may have given you complete freedom, provided only that you attribute them as the owner of the work. There are a few simple choices and a limited menu of permutations.

What makes these licenses unusual is that they can be read by two groups that normal licenses exclude—human beings (rather than just lawyers) and computers. The textbooks, photos, films, and songs have a tasteful little emblem on them marked with a “cc” which, if you click on it, links to a “Commons Deed,” a simple one-page explanation of the freedoms you have. There are even icons—a dollar with a slash through it, for example—that make things even clearer. Better still, the reason the search engines could find this material is that the licenses also “tell” search engines exactly what freedoms have been given. Simple “metadata” (a fancy word for tags that computers can read) mark the material with its particular level of freedoms. This is not digital rights management. The license will not try to control your computer, install itself on your hard drive, or break your TV. It is just an expression of the terms under which the author has chosen to release the work. That means that if you search Google or Flickr for “works I am free to share, even commercially,” you know you can go into business selling those textbooks, or printing those photos on mugs and T-shirts, so long as you give the author attribution. If you search for “show me works I can build on,” you know you are allowed to make what copyright lawyers call “derivative works.”

The idea behind Creative Commons was simple. As I pointed out in the first chapter, copyright adheres automatically on “fixation.” As soon as you lift the pen from the paper, click the shutter, or save the file, the work is copyrighted. No formalities. No need even to use the little symbol ©. Once copyrighted, the work is protected by the full might of the legal system. And the legal system’s default setting is that “all rights are reserved” to the author, which means effectively that anyone but the author is forbidden to copy, adapt, or publicly perform the work. This might have been a fine rule for a world in which there were high barriers to publication. The material that was not published was theoretically under an “all rights reserved” regime, but who cared?

It was practically inaccessible anyway. After the development of the World Wide Web, all that had changed. Suddenly people and institutions, millions upon millions of them, were putting content online—blogs, photo sites, vlogs, podcasts, course materials. It was all just up there.

But what could you do with it? You could read it, or look at it, or play it presumably—otherwise why had the author put it up? But could you copy it? Put it on your own site? Include it in a manual used by the whole school district? E-mail it to someone? Translate it into your own language? Quote beyond the boundaries of fair use? Adapt for your own purposes? Take the song and use it for your video? Of course, if you really wanted the work a lot, you could try to contact the author—not always easy. And one by one, we could all contact each other and ask for particular types of permissions for use. If the use was large enough or widespread enough, perhaps we would even think that an individual contract was necessary. Lawyers could be hired and terms hashed out.

All this would be fine if the author wished to retain all the rights that copyright gives and grant them only individually, for pay, with lawyers in the room. But what about the authors, the millions upon millions of writers, and photographers and musicians, and filmmakers and bloggers and scholars, who very much *want* to share their work? The Cora Beth Bridges of the world are never going to write individual letters to the Colin Mutchlers of the world asking for permission to make a derivative work out of “My Life.” The person who translated my articles into Spanish or Mandarin, or the people who repost them on their Web sites, or include them in their anthologies might have asked permission if I had not granted it in advance. I doubt though that I would have been contacted by the very talented person who took images from a comic book about fair use that I co-wrote and mashed them up with words from a book by Larry Lessig, and some really nice music from someone none of us had ever met. Without some easy way to give permission in advance, and to do so in a way that human beings and computers, as well as lawyers, can understand, those collaborations will never happen, though all the parties would be delighted if they did. These are losses from “failed sharing”—every bit as real as losses from unauthorized copying, but much less in the public eye.

Creative Commons was conceived as a private “hack” to produce a more fine-tuned copyright structure, to replace “all rights reserved” with “some rights reserved” for those who wished to do so. It tried to do for culture what the General Public License had done for software. It made use of the same technologies that had created the issue: the technologies that made fixation of

expressive content and its distribution to the world something that people, as well as large concentrations of capital, could do. As a result, it was able to attract a surprising range of support—Jack Valenti of the Motion Picture Association of America and Hillary Rosen of the Recording Industry Association of America, as well as John Perry Barlow of the Grateful Dead, whose attitude toward intellectual property was distinctly less favorable. Why could they all agree? These licenses were not a choice forced on anyone. The author was choosing what to share and under what terms. But that sharing created something different, something new. It was more than a series of isolated actions. The result was the creation of a global “commons” of material that was open to all, provided they adhered to the terms of the licenses. Suddenly it was possible to think of creating a work entirely out of Creative Commons-licensed content—text, photos, movies, music. Your coursebook on music theory, or your documentary on the New York skyline, could combine your own original material with high-quality text, illustrations, photos, video, and music created by strangers. One could imagine entire fields—of open educational content or of open music—in which creators could work without keeping one eye nervously on legal threats or permissions.

From one perspective, Creative Commons looks like a simple device for enabling exercise of authorial control, remarkable only for the extremely large number of authors making that choice and the simplicity with which they can do so. From another, it can be seen as re-creating, by private choice and automated licenses, the world of creativity before law had permeated to the finest, most atomic level of science and culture—the world of folk music or 1950s jazz, of jokes and slang and recipes, of Ray Charles’s “rewording” of gospel songs, or of Isaac Newton describing himself as “standing on the shoulders of giants” (and not having to pay them royalties). Remember, that is not a world *without* intellectual property. The cookbook might be copyrighted even if the recipe was not. Folk music makes it to the popular scene and is sold as a copyrighted product. The jazz musician “freezes” a particular version of the improvisation on a communally shared set of musical motifs, records it, and sometimes even claims ownership of it. Newton himself was famously touchy about precedence and attribution, even if not about legal ownership of his ideas. But it is a world in which creativity and innovation proceed on the basis of an extremely large “commons” of material into which it was never imagined that property rights could permeate.

For many of us, Creative Commons was conceived of as a second-best solution created by private agreement because the best solution could not be

obtained through public law. The best solution would be a return of the formality requirement—a requirement that one at least write the words “James Boyle copyright 2008,” for example, in order to get more than 100 years of legal protection backed by “strict liability” and federal criminal law. Those who did not wish to have the legal monopoly could omit the phrase and the work would pass into the public domain, with a period of time during which the author could claim copyright retrospectively if the phrase was omitted by accident. The default position would become freedom and the dead weight losses caused by giving legal monopolies to those who had not asked for them, and did not want them, would disappear. To return to the words of Justice Brandeis that I quoted at the beginning of the book:

The general rule of law is, that the noblest of human productions—knowledge, truths ascertained, conceptions, and ideas—become, after voluntary communication to others, free as the air to common use. Upon these incorporeal productions the attribute of property is continued after such communication only in certain classes of cases where public policy has seemed to demand it.

Brandeis echoes the Jeffersonian preference for a norm of freedom, with narrowly constrained exceptions only when necessary. That preference means that the commons of which I spoke is a relatively large one—property rights are the exception, not the norm. Of course, many of those who use Creative Commons licenses might disagree with that policy preference and with every idea in this book. They may worship the DMCA or just want a way to get their song or their article out there while retaining some measure of control. That does not matter. The licenses are agnostic. Like a land trust which has a local pro-growth industrialist and a local environmentalist on its board, they permit us to come to a restricted agreement on goals (“make sure this space is available to the public”) even when underlying ideologies differ. They do this using those most conservative of tools—property rights and licenses. And yet, if our vision of property is “sole and despotic dominion,” these licenses have created something very different—a commons has been made out of private and exclusive rights.

My point here is that Creative Commons licenses or the tools of free and open source software—to which I will turn in a moment—represent something more than merely a second-best solution to a poorly chosen rule. They represent a visible example of a type of creativity, of innovation, which has been around for a very long time, but which has reached new salience on the Internet—distributed creativity based around a shared commons of material.

FREE AND OPEN SOURCE SOFTWARE

In 2007, Clay Shirky, an incisive commentator on networked culture, gave a speech which anyone but a Net aficionado might have found simultaneously romantic and impenetrable. He started by telling the story of a Shinto shrine that has been painstakingly rebuilt to exactly the same plan many times over its 1,300-year life—and which was denied certification as a historic building as a result. Shirky’s point? What was remarkable was not the building. It was a community that would continue to build and rebuild the thing for more than a millennium.

From there, Shirky shifted to a discussion of his attempt to get AT&T to adopt the high-level programming language Perl—which is released as free and open source software under the General Public License. From its initial creation by Larry Wall in 1987, Perl has been adapted, modified, and developed by an extraordinary range of talented programmers, becoming more powerful and flexible in the process. As Shirky recounts the story, when the AT&T representatives asked “where do you get your support?” Shirky responded, “‘we get our support from a community’—which to them sounded a bit like ‘we get our Thursdays from a banana.’” Shirky concluded the speech thus:

We have always loved one another. We’re human. It’s something we’re good at. But up until recently, the radius and half-life of that affection has been quite limited. With love alone, you can plan a birthday party. Add coordinating tools and you can write an operating system. In the past, we would do little things for love, but big things required money. Now we can do big things for love.¹

There are a few people out there for whom “operating systems” and “love” could plausibly coexist in a sentence not constructed by an infinite number of monkeys. For most though, the question is, what could he possibly have meant?

The arguments in this book so far have taken as a given the incentives and collective action problems to which intellectual property is a response. Think of Chapter 1 and the economic explanation of “public goods.” The fact that it is expensive to do the research to find the right drug, but cheap to manufacture it once it is identified provides a reason to create a legal right of exclusion. In those realms where the innovation would not have happened anyway, the legal right of exclusion gives a power to price above cost, which in turn gives incentives to creators and distributors. So goes the theory. I have discussed the extent to which the logic of enclosure works for the commons of the mind as well as it did for the arable commons, taking into account the effects of an information

society and a global Internet. What I have not done is asked whether a global network actually transforms some of our assumptions about how creation happens in a way that reshapes the debate about the need for incentives, at least in certain areas. This, however, is exactly the question that needs to be asked.

For anyone interested in the way that networks can enable new collaborative methods of production, the free software movement, and the broader but less political movement that goes under the name of open source software, provide interesting case studies.² Open source software is released under a series of licenses, the most important being the General Public License (GPL). The GPL specifies that anyone may copy the software, provided the license remains attached and the source code for the software always remains available.³ Users may add to or modify the code, may build on it and incorporate it into their own work, but if they do so, then the new program created is also covered by the GPL. Some people refer to this as the “viral” nature of the license; others find the term offensive.⁴ The point, however, is that the open quality of the creative enterprise spreads. It is not simply a donation of a program or a work to the public domain, but a continual accretion in which all gain the benefits of the program on pain of agreeing to give their additions and innovations back to the communal project.

For the whole structure to work without large-scale centralized coordination, the creation process has to be modular, with units of different sizes and complexities, each requiring slightly different expertise, all of which can be added together to make a grand whole. I can work on the sendmail program, you on the search algorithms. More likely, lots of people try, their efforts are judged by the community, and the best ones are adopted. Under these conditions, this curious mix of Kropotkin and Adam Smith, Richard Dawkins and Richard Stallman, we get distributed production without having to rely on the proprietary exclusion model. The whole enterprise will be much, much, much greater than the sum of the parts.

What’s more, and this is a truly fascinating twist, when the production process does need more centralized coordination, some governance that guides how the sticky modular bits are put together, it is at least theoretically possible that we can come up with the control system *in exactly the same way*. In this sense, distributed production is potentially recursive. Governance processes, too, can be assembled through distributed methods on a global network, by people with widely varying motivations, skills, and reserve prices.⁵

The free and open source software movements have produced software that rivals or, some claim, exceeds the capabilities of conventional proprietary,

binary-only software.⁶ Its adoption on the “enterprise level” is impressive, as is the number and enthusiasm of the various technical testaments to its strengths. You have almost certainly used open source software or been its beneficiary. Your favorite Web site or search engine may run on it. If your browser is Firefox, you use it every day. It powers surprising things around you—your ATM or your TiVo. The plane you are flying in may be running it. It just works.

Governments have taken notice. The United Kingdom, for example, concluded last year that open source software “will be considered alongside proprietary software and contracts will be awarded on a value-for-money basis.” The Office of Government Commerce said open source software is “a viable desktop alternative for the majority of government users” and “can generate significant savings. . . . These trials have proved that open source software is now a real contender alongside proprietary solutions. If commercial companies and other governments are taking it seriously, then so must we.”⁷ Sweden found open source software to be in many cases “equivalent to—or better than—commercial products” and concluded that software procurement “shall evaluate open software as well as commercial solutions, to provide better competition in the market.”⁸

What is remarkable is not merely that the software works technically, but that it is an example of widespread, continued, high-quality innovation. The really remarkable thing is that it works socially, as a continuing system, sustained by a network consisting both of volunteers and of individuals employed by companies such as IBM and Google whose software “output” is nevertheless released into the commons.

Here, it seems, we have a classic public good: code that can be copied freely and sold or redistributed without paying the creator or creators. This sounds like a tragedy of the commons of the kind that I described in the first three chapters of the book. Obviously, with a nonrival, nonexcludable good like software, this method of production cannot be sustained; there are inadequate incentives to ensure continued production. *E pur si muove*, as Galileo is apocryphally supposed to have said in the face of Cardinal Bellarmine’s certainties: “And yet it moves.”⁹ Or, as Clay Shirky put it, “we get our support from a community.”

For a fair amount of time, most economists looked at open source software and threw up their hands. From their point of view, “we get our support from a community” did indeed sound like “we get our Thursdays from a banana.” There is an old economics joke about the impossibility of finding a twenty-dollar bill lying on a sidewalk. In an efficient market, the money would already have been picked up. (Do not wait for a punch line.) When economists

looked at open source software they saw not a single twenty-dollar bill lying implausibly on the sidewalk, but whole bushels of them. Why would anyone work on a project the fruits of which could be appropriated by anyone? Since copyright adheres on fixation—since the computer programmer already has the legal power to exclude others—why would he or she choose to take the extra step of adopting a license that undermined that exclusion? Why would anyone choose to allow others to use and modify the results of their hard work? Why would they care whether the newcomers, in turn, released their contributions back into the commons?

The puzzles went beyond the motivations of the people engaging in this particular form of “distributed creativity.” How could these implausible contributions be organized? How should we understand this strange form of organization? It is not a company or a government bureaucracy. What could it be? To Richard Epstein, the answer was obvious and pointed to a reason the experiment must inevitably end in failure:

The open source movement shares many features with a workers’ commune, and is likely to fail for the same reason: it cannot scale up to meet its own successes. To see the long-term difficulty, imagine a commune entirely owned by its original workers who share pro rata in its increases in value. The system might work well in the early days when the workforce remains fixed. But what happens when a given worker wants to quit? Does that worker receive in cash or kind his share of the gain in value during the period of his employment? If not, then the run-up in value during his period of employment will be gobbled up by his successor—a recipe for immense resentment. Yet that danger can be ducked only by creating a capital structure that gives present employees separable interests in either debt or equity in exchange for their contributions to the company. But once that is done, then the worker commune is converted into a traditional company whose shareholders and creditors contain a large fraction of its present and former employers. The bottom line is that idealistic communes cannot last for the long haul.¹⁰

There are a number of ideas here. First, “idealistic communes cannot last for the long haul.” The skepticism about the staying power of idealism sounds plausible today, though there are some relatively prominent counterexamples. The Catholic Church is also a purportedly idealistic institution. It is based on canonical texts that are subject to even more heated arguments about textual interpretation than those which surround the General Public License. It seems to be surviving the long haul quite well.

The second reason for doomsaying is provided by the word “commune.” The problems Epstein describes are real where tangible property and excludable

assets are involved. But is the free and open source community a “commune,” holding tangible property in common and excluding the rest of us? Must it worry about how to split up the proceeds if someone leaves because of bad karma? Or is it a community creating and offering to the world the ability to use, for free, nonrival goods that all of us can have, use, and reinterpret as we wish? In that kind of commune, each of us could take all the property the community had created with us when we left and the commune would still be none the poorer. Jefferson was not thinking of software when he talked of the person who lights his taper from mine but does not darken me, but the idea is the same one. Copying software is not like fighting over who owns the scented candles or the VW bus. Does the person who wrote the “kernel” of the operating system resent the person who, much later, writes the code to manage Internet Protocol addresses on a wireless network? Why should he? Now the program does more cool stuff. Both of them can use it. What’s to resent?

How about idealism? There is indeed a broad debate on the reasons that the system works: Are the motivations those of the gift economy? Is it, as Shirky says, simply the flowering of an innate love that human beings have always had for each other and for sharing, now given new strength by the geographic reach and cooperative techniques the Internet provides? “With love alone, you can plan a birthday party. Add coordinating tools and you can write an operating system.” Is this actually a form of potlatch, in which one gains prestige by the extravagance of the resources one “wastes”? Is open source an implicit résumé-builder that pays off in other ways? Is it driven by the species-being, the innate human love of creation that continually drives us to create new things even when *homo economicus* would be at home in bed, mumbling about public goods problems?¹¹

Yochai Benkler and I would argue that these questions are fun to debate but ultimately irrelevant.¹² Assume a random distribution of incentive structures in different people, a global network—transmission, information sharing, and copying costs that approach zero—and a modular creation process. With these assumptions, it just does not matter why they do it. In lots of cases, they *will* do it. One person works for love of the species, another in the hope of a better job, a third for the joy of solving puzzles, and a fourth because he has to solve a particular problem anyway for his own job and loses nothing by making his hack available for all. Each person has their own reserve price, the point at which they say, “Now I will turn off *Survivor* and go and create something.” But on a global network, there are a lot of people, and with numbers that big and information overhead that small, even relatively hard projects

will attract motivated and skilled people whose particular reserve price has been crossed.

More conventionally, many people write free software because they are paid to do so. Amazingly, IBM now earns more from what it calls “Linux-related revenues” than it does from traditional patent licensing, and IBM is the largest patent holder in the world.¹³ It has decided that the availability of an open platform, to which many firms and individuals contribute, will actually allow it to sell more of its services, and, for that matter, its hardware. A large group of other companies seem to agree. They like the idea of basing their services, hardware, and added value on a widely adopted “commons.” This does not seem like a community in decline.

People used to say that collaborative creation could never produce a quality product. That has been shown to be false. So now they say that collaborative creation cannot be sustained because the governance mechanisms will not survive the success of the project. Professor Epstein conjures up a “central committee” from which insiders will be unable to cash out—a nice mixture of communist and capitalist metaphors. All governance systems—including democracies and corporate boards—have problems. But so far as we can tell, those who are influential in the free software and open source governance communities (there is, alas, no “central committee”) feel that they are doing very well indeed. In the last resort, when they disagree with decisions that are taken, there is always the possibility of “forking the code,” introducing a change to the software that not everyone agrees with, and then letting free choice and market selection converge on the preferred iteration. The free software ecosystem also exhibits diversity. Systems based on GNU-Linux, for example, have distinct “flavors” with names like Ubuntu, Debian, and Slackware, each with passionate adherents and each optimized for a particular concern—beauty, ease of use, technical manipulability. So far, the tradition of “rough consensus and running code” seems to be proving itself empirically as a robust governance system.

Why on earth should we care? People have come up with a surprising way to create software. So what? There are at least three reasons we might care. First, it teaches us something about the limitations of conventional economics and the counterintuitive business methods that thrive on networks. Second, it might offer a new tool in our attempt to solve a variety of social problems. Third, and most speculative, it hints at the way that a global communications network can sometimes help move the line between work and play, professional and amateur, individual and community creation, rote production and compensated “hobby.”

We should pay attention to open source software because it shows us something about business methods in the digital world—indeed in the entire world of “information-based” products, which is coming to include biotechnology. The scale of your network matters. The larger the number of people who use your operating system, make programs for your type of computer, create new levels for your game, or use your device, the better off you are. A single fax machine is a paperweight. Two make up a communications link. Ten million and you have a ubiquitous communications network into which your “paperweight” is now a hugely valuable doorway.

This is the strange characteristic of networked goods. The actions of strangers dramatically increase or decrease the usefulness of your good. At each stage the decision of someone else to buy a fax machine increases the value of mine. If I am eating an apple, I am indifferent about whether you are too. But if I have a fax machine then my welfare is actually improved by the decisions of strangers to buy one. The same process works in reverse. Buy a word processing program that becomes unpopular, get “locked in” to using it, and find yourself unable to exchange your work easily with others. Networks matter and increasing the size of the networks continues to add benefits to the individual members.

What’s true for the users of networks is doubly so for the producers of the goods that create them. From the perspective of a producer of a good that shows strong network effects such as a word processing program or an operating system, the optimal position is to be the company that owns and controls the dominant product on the market. The ownership and control is probably by means of intellectual property rights, which are, after all, the type of property rights one finds on networks. The value of that property depends on those positive and negative network effects. This is the reason Microsoft is worth so much money. The immense investment in time, familiarity, legacy documents, and training that Windows or Word users have provides a strong incentive not to change products. The fact that other users are similarly constrained makes it difficult to manage any change. Even if I change word processor formats and go through the trouble to convert all my documents, I still need to exchange files with you, who are similarly constrained. From a monopolist’s point of view, the handcuffs of network effects are indeed golden, though opinions differ about whether or not this is a cause for antitrust action.

But if the position that yields the most revenue is that of a monopolist exercising total control, the second-best position may well be that of a company contributing to a large and widely used network based on open standards and,

perhaps, open software. The companies that contribute to open source do not have the ability to exercise monopoly control, the right to extract every last cent of value from it. But they do have a different advantage; they get the benefit of all the contributions to the system without having to pay for them. The person who improves an open source program may not work for IBM or Red Hat, but those companies benefit from her addition, just as she does from theirs. The system is designed to continue growing, adding more contributions back into the commons. The users get the benefit of an ever-enlarging network, while the openness of the material diminishes the lock-in effects. Lacking the ability to extract payment for the network good itself—the operating system, say—the companies that participate typically get paid for providing tied goods and services, the value of which increases as the network does.

I write a column for the *Financial Times*, but I lack the fervor of the true enthusiast in the “Great Game of Markets.” By themselves, counterintuitive business methods do not make my antennae tingle. But as Larry Lessig and Yochai Benkler have argued, this is something more than just another business method. They point us to the dramatic role that openness—whether in network architecture, software, or content—has had in the success of the Internet. *What is going on here is actually a remarkable corrective to the simplistic notion of the tragedy of the commons, a corrective to the Internet Threat storyline and to the dynamics of the second enclosure movement.* This commons creates and sustains value, and allows firms and individuals to benefit from it, without depleting the value already created. To appropriate a phrase from Carol Rose, open source teaches us about the *comedy* of the commons, a way of arranging markets and production that we, with our experience rooted in physical property and its typical characteristics, at first find counterintuitive and bizarre. Which brings us to the next question for open source. Can we use its techniques to solve problems beyond the world of software production?

In the language of computer programmers, the issue here is “does it scale?” Can we generalize anything from this limited example? How many types of production, innovation, and research fit into the model I have just described? After all, for many innovations and inventions one needs hardware, capital investment, and large-scale, real-world data collection—*stuff*, in its infinite recalcitrance and facticity. Maybe the open source model provides a workaround to the individual incentives problem, but that is not the only problem. And how many types of innovation or cultural production are as

modular as software? Is open source software a paradigm case of collective innovation that helps us to understand open source software and not much else?

Again, I think this is a good question, but it may be the wrong one. My own guess is that an open source method of production is far more common than we realize. “Even before the Internet” (as some of my students have taken to saying portentously), science, law, education, and musical genres all developed in ways that are markedly similar to the model I have described. The marketplace of ideas, the continuous roiling development in thought and norms that our political culture spawns, owes much more to the distributed, nonproprietary model than it does to the special case of commodified innovation that we think about in copyright and patent. Not that copyright and patent are unimportant in the process, but they may well be the exception rather than the norm. Commons-based production of ideas is hardly unfamiliar, after all.

In fact, all the mottos of free software development have their counterparts in the theory of democracy and open society; “given enough eyeballs, all bugs are shallow” is merely the most obvious example. Karl Popper would have cheered.¹⁴ The importance of open source software is not that it introduces us to a wholly new idea. It is that it makes us see clearly a very old idea. With open source the technology was novel, the production process transparent, and the result of that process was a “product” which outcompeted other products in the marketplace. “How can this have happened? What about the tragedy of the commons?” we asked in puzzlement, coming only slowly to the realization that other examples of commons-based, nonproprietary production were all around us.

Still, this does not answer the question of whether the model can scale still further, whether it can be applied to solve problems in other spheres. To answer that question we would need to think more about the modularity of other types of inventions. How much can they be broken down into chunks suitable for distribution among a widespread community? Which forms of innovation have some irreducible need for high capital investment in distinctly nonvirtual components—a particle accelerator or a Phase III drug trial? Again, my guess is that the increasing migration of the sciences toward data- and processing-rich models makes much more of innovation and discovery a potential candidate for the distributed model. Bioinformatics and computational biology, the open source genomics project,¹⁵ the BioBricks Foundation I mentioned in the last chapter, the possibility of distributed data

scrutiny by lay volunteers¹⁶—all of these offer intriguing glances into the potential for the future. Finally, of course, the Internet is one big experiment in, as Benkler puts it, peer-to-peer cultural production.¹⁷

If these questions are good ones, why are they also the wrong ones? I have given my guesses about the future of the distributed model of innovation. My own utopia has it flourishing alongside a scaled-down, but still powerful, intellectual property regime. Equally plausible scenarios see it as a dead end or as the inevitable victor in the war of productive processes. These are all guesses, however. At the very least, there is some possibility, even hope, that we could have a world in which much more of intellectual and inventive production is free. “‘Free’ as in ‘free speech,’” Richard Stallman says, not “free as in ‘free beer.’”¹⁸ But we could hope that much of it would be both free of centralized control and low- or no-cost. When the marginal cost of reproduction is zero, the marginal cost of transmission and storage approaches zero, the process of creation is additive, and much of the labor doesn’t charge, the world looks a little different.¹⁹ This is at least a possible future, or part of a possible future, and one that we should not foreclose without thinking twice. Yet that is what we are doing. The Database Protection Bills and Directives, which extend intellectual property rights to the layer of facts;²⁰ the efflorescence of software patents;²¹ the UCITA-led validation of shrinkwrap licenses that bind third parties;²² the Digital Millennium Copyright Act’s anticircumvention provisions²³—the point of all of these developments is not merely that they make the peer-to-peer model difficult, but that in many cases they rule it out altogether. I will assert this point here, rather than argue for it, but I think it can be (and has been) demonstrated quite convincingly.²⁴

The point is, then, that there is a chance that a new (or old, but under-recognized) method of production could flourish in ways that seem truly valuable—valuable to free speech, innovation, scientific discovery, the wallets of consumers, to what William Fisher calls “semiotic democracy,”²⁵ and, perhaps, valuable to the balance between joyful creation and drudgery for hire. True, it is only a chance. True, this theory’s scope of operation and sustainability are uncertain. But why would we want to foreclose it? That is what the recent expansions of intellectual property threaten to do. And remember, these expansions were dubious even in a world where we saw little or no possibility of the distributed production model I have described, where discussion of network effects had yet to reach the pages of *The New Yorker*,²⁶ and where our concerns about the excesses of intellectual property were simply the ones that Jefferson, Madison, and Macaulay gave us so long ago.

LEARNING FROM THE SHARING ECONOMY

Accept for the sake of argument that the free software community actually works, actually produces high-quality products capable of competing in the market with proprietary alternatives. Concede for a moment that the adoption of Creative Commons licenses shows there are millions of creators out there who want to share their works with others. Many of those creators even want to allow the world to build on their material. Indeed, let us concede that the whole history of the Web, from Wikipedia to the obsessive and usefully detailed sites created on everything from Vikings to shoe polishes, shows a desire to share one's knowledge, to build on the work of others one has never met. These efforts are remarkably varied. Some are ultimately aimed at profit—even if their results are free. Think of IBM's open source initiatives or musicians who release Creative Commons-licensed work in order to get more club gigs. Some are provided as a volunteer act of benevolence or civic duty, even if they “compete” with expensive proprietary alternatives. Think of Wikipedia or MIT's OpenCourseWare. When the infrastructure for this collaboration does not exist, it gets assembled—and quickly. Both the GPL and Creative Commons are examples. Accept all of this. So what?

Lesson number one comes from the nonprofit activities—everything from Wikipedia to Web sites created by enthusiasts. People like to create and wish to share. In many cases they will do so without financial reward. A surprising amount of useful, creative, or expressive activity is generated without any financial incentive at all.

Should this cause us to throw out the economic case for copyrights? No. But it should lead us to reassess it. As I explained in Chapter 1, copyright provides an incentive for two distinct activities. First, it offers an incentive to create the work in the first place. The author of *Windows for Dummies* or *Harry Potter* gets a right to exclude others from copying the work, a right that he or she can sell in the marketplace. The goal is to offer a financial reason to devote time to this particular creative activity. It is this incentive that is most often cited when attempting to persuade policy makers to expand protection. Second, it offers an incentive to distribute the work—to typeset and print large quantities of the work and to sell it to bookstores, or to broadcast it, or put it on movie screens.

Each medium is economically different, of course. The economics of the feature film are different from those of the book, the magazine, or the operating system. Thus, we have never had very good figures on the relative importance of

these incentives. We can only guess at how much of the incentive from copyright goes to encouraging creation and how much to distribution. Until recently, most types of distribution demanded higher levels of capital. The industry structure that resulted often consisted of creators who worked as wage or contract labor for distributors—either never acquiring copyright in their work in the first place or immediately transferring that copyright to their employers. Because distribution was expensive, our experience with material generated for fun or out of a love of sharing was an essentially private and local one. You might have a neighbor's photocopied sheet of baking recipes that worked well at high altitudes, or of fishing techniques that worked well on a particular lake, a song that a friend created for a special occasion, or a short story you wrote for your kids—and then typed up for them to tell to theirs. Financial incentives were not needed to encourage the creation of the work, but the cost of distribution dramatically limited its dissemination.

The single most dramatic thing that the Web has done by lowering the cost of communication and distribution, at the same moment that other electronic tools lowered the cost of production, is to make this local and private activity a global and public one. Someone, somewhere, will have written the guide to fishing on that lake, baking at that altitude, washing windows, or treating stings from Portuguese man-of-war jellyfish. Someone will have taken a photo of the Duke Chapel or explained the history, economics, and chemistry of shoe polish or distilling. Someone might even have created a great class on music theory or C++ programming. Someone will have written a handy little program to manage DNS requests on a local network. Bizarrely, at least as far as the economists were concerned, these people all wanted to share what they had made. Because of the genius of search engines, and the implicit peer-review function that those engines deduce from patterns of links to pages, I can find that material when I need it.

True, much of the material on the Web is inane or insane, confused, badly written, tendentious, and inaccurate. (It should be noted that this is hardly a problem confined to the Web or volunteer-generated material. Personally, I would not want *People* magazine or Fox News in a time capsule to represent my civilization. But some of the material on the Web is clearly worse.) Yes, Wikipedia is occasionally inaccurate—though in one test in *Nature* it stacked up well against the *Encyclopedia Britannica*, and it is obviously much more encyclopedic in its coverage. But all of this misses the point.

Consider how your expectations about information retrieval have changed in the last fifteen years. We now simply assume that questions about a piece of

architecture, a bit of local history, a recipe, or the true author of a song can all be answered within seconds. We have forgotten what it is like to be routinely in ignorance because of the unavailability of some piece of information. One podcaster I talked to called it being a member of “the right-click generation”: “When I am walking around and I see a building, I almost feel as though I ought to be able to ‘right click’ it and have the architect’s name pop up.” Consider that it now seems normal for a gay Iraqi man in Baghdad to have a blog that offers hundreds of thousands of readers around the world a literate and touching account of the American occupation from a perspective entirely different from that provided by the mainstream press.²⁷ We think it normal for a person of moderate resources to be able to speak to the world from a war zone, whether or not he is affiliated with a newspaper or credentialed by a corporation.

These examples are not the end of the process. Our methods of sorting, ranking, and verifying the material generated are still evolving. They may improve even beyond this point. We are only fifteen years into this particular experiment, after all. And a huge amount of this material is produced by our fellow citizens without the profit motive.

Does this mean that we no longer need copyright or patent protection to encourage the production and distribution of creative work? No. The fishing tips are great, but I still might buy a handsomely illustrated guide to take on the lake with me or, even better, just stay at home and read *A River Runs Through It*. *The New Yorker*, and not a sheaf of printouts from the Web, still sits on my coffee table, though much of the high-quality content I read comes to me online, for free, from strangers who are generating it for pleasure, not profit, or who profit from open sharing, not closed control. The online blogosphere provides a vital counterpoint to mainstream media, but it exists in a symbiotic—some would say parasitic—relationship with that media and the network of professional news gatherers for which it pays. Some of the most interesting open source production methods actually rely on copyright. Even if they did not, open source production would not suffice to run our pharmaceutical industry (though it might help with certain stages of the drug discovery process).

Still, just as it would be silly to dismiss the importance of intellectual property based on our experience of blogs and Wikipedia and open source software, it would be equally silly to underestimate what the Web has taught us. The Web has enabled an astonishing flowering of communication and expression, an astounding democratization of creativity. We have learned just how

strong, and how useful, is the human urge to express, communicate, invent, and create—provided the barriers to sharing are lowered. These are the very things that copyright and patent are supposed to encourage. For us to portray the Web—as the Internet Threat story line does—as predominantly a *threat* to creativity is simply perverse. For us to base our policies only on that notion would be a tragedy. We might end up stultifying one of the greatest explosions of human creativity the world has ever seen by treating it as an unimportant marginal case and instead designing our rules around the production processes of commercial culture in the late twentieth century.

The shape of our copyright and to a lesser extent our patent system comes from a world in which almost all large-scale distribution was an expensive, capital-intensive enterprise. The roles of gatekeeper and financier, producer and assembler, distributor and advertiser, tended naturally to coalesce into vertically integrated firms or symbiotic commercial partnerships. Those firms were presumed to be the proxy for the public interest when it came to intellectual property policy. Who would know better than they what was needed? Occasionally, device manufacturers would provide a counterweight—as in the *Sony* case—where the defense of a particular “consumer freedom” actually created a market for a complementary product. Artists and authors might be trotted out as appealing spokespersons, though the laws that were made only sporadically reflected their economic and artistic interests. Librarians and educational institutions had influence at the edges. Most of the time, though, it was the assemblers and distributors of content whose voices and assumptions about markets would be heard.

Out of this pattern of habit and influence, and out of much deeper notions about authorship and invention that I have explored elsewhere, developed an ideology, a worldview. Call it maximalism. Its proponents sincerely believed in it and pursued it even when it did not make economic sense. (Think how lucky the movie industry is that it lost the *Sony* case.) It has been the subject of this book. Its tenets are that intellectual property is just like physical property, that rights need to increase proportionately as copying costs decrease, and that, in general, increasing levels of intellectual property protection will yield increasing levels of innovation. Despite its defense of ever-increasing government-granted monopolies, this ideology cloaks itself in the rhetoric of free markets. The bumbling state, whose interventions in the economy normally spell disaster, turns into a scalpel-wielding genius when its monopolies and subsidies are provided through intellectual property rights rather than regulatory fiat. Above all, this way of seeing the world minimizes the impor-

tance of creativity, expression, and distribution that takes place outside its framework and ignores or plays down the importance of the input side of the equation—the need to focus on the material from which culture and science are made, as well as the protected expression and inventions made from that raw material.

This process was not—let me stress—was *not* a simple process of economic determinism or industry conspiracy. Anyone who claims that is the thesis of this book simply has not read it. (Reviewers beware.) Let us start with economic determinism. It was not a situation in which the law mechanically recorded the interests of the most economically important industries in the area. This was the creation of a worldview, not the steely-eyed calculation of profit and loss. Not only did many of the rules we ended up with make no sense from the point of view of some of the largest economic players in the area—think of the device manufacturers, the search engines, and so on—they frequently made no sense from the perspective of those proposing them. Attempting to twist the law to make it illegal for technology to interfere with your old business method is frequently bad for the industry seeking the protection, as well as for the technology, the market, and the wider society. Since this worldview makes incumbents systematically blind to profit-making opportunities that could be secured by greater openness, rather than greater control, it actually disables them from pursuing some of the most promising methods by which they could have made money for their shareholders. Again, the chapter on the *Sony* decision offers a salutary example.

Economic determinism does not explain the rules we have. Neither are those rules simply a result of the manipulation of elected officials by incumbent industries through crafty campaign contributions and distorted evidence (though to be sure, there was a lot of that as well). Many of the people who put forward this worldview—both lobbyists and lobbied—sincerely believe that more rights will always lead to more innovation, that all property rights are the same, that we do not need to think about both the input and output sides of the equation, that cheaper copying techniques automatically require greater protections, and so on.

What of the modest suggestions I put forward here? We could sum them up thus: do not apply identical assumptions to physical and intellectual property. Focus on both the inputs to and the outputs of the creative process; protecting the latter may increase the cost of the former. Look both at the role of the public domain and the commons of cultural and scientific material *and* at the need to provide incentives for creativity and distribution through exclusive

rights. More rights will not automatically produce more innovation. Indeed, we should confine rights as narrowly as possible while still providing the desired result. Look at the empirical evidence before and after increasing the level of protection. Pay attention to the benefits as well as the costs of the new technologies and the flowering of creativity they enable.

To me, these points seem bland, boring, obvious—verging on tautology or pabulum. To many believers in the worldview I have described, they are either straightforward heresy or a smokescreen for some real, underlying agenda—which is identified as communism, anarchism, or, somewhat confusingly, both.

This account smacks of exaggeration, I know. How could things be so one-sided? The best answer I can give came from a question I was asked at a recent conference. The questioner pointed out politely that it was unlikely that the policy-making process would ignore such a fundamental and obvious set of points—points that I myself observed had been well understood for hundreds of years. I had used many examples of intellectual property rights being extended—in length, breadth, scope. Why had I not spoken, he asked, of all the times over the last fifty years when intellectual property rights had been weakened, curtailed, shortened? Since human beings were fallible, surely there were occasions when the length of a copyright or patent term had proved to be too long, or the scope of a right too large, and the rights had been narrowed appropriately by legislation. Why did I not cite any of these? The answer is simple. To the best of my knowledge, there are none. Legislatively, intellectual property rights have moved only in one direction—outward. (Court decisions present a more complex picture, as the previous chapter's discussion of software copyrights and business method patents shows.)

What are the odds that the costs of new technologies are *always* greater than their benefits as far as intellectual property rights holders are concerned? This pattern is not a matter of policies carefully crafted around the evidence. It is the fossil record of fifty years of maximalism. If I lean toward the other side of the story it is not because I am a foe of intellectual property. It is because I believe our policies have become fundamentally unbalanced—unbalanced in ways that actually blind us to what is going on in the world of creativity.

We are living through an existence-proof that there are other methods of generating innovation, expression, and creativity than the proprietary, exclusionary model of sole control. True, these methods existed before. Yet they tended to be local or invisible or both. The Internet has shown conclusively and visibly that—at least in certain sectors—we can have a global flowering of

creativity, innovation, and information sharing in which intellectual property rights function in a very different way than under the standard model of proprietary control. In some cases, intellectual property rights were simply irrelevant—much of the information sharing and indexing on the Web falls within this category. In some cases they were used to *prevent* exclusivity. Think of Creative Commons or the General Public License. In some, they were actually impediments. Software patents, for example, have a negative effect on open source software development—one that policy makers are only now slowly beginning to acknowledge.

It is important not to overstate how far the sharing economy can get us. It might help to cut the costs of early-stage drug development, as the Tropical Disease Initiative attempts to do for neglected diseases. It will not generate a Phase III drug trial or bring a drug to market. Sharing methods might be used to generate cult movies such as *Star Wreck: In the Pirkinning*, which was created using techniques borrowed from open source software and is available under a Creative Commons license. They will not produce a mammoth blockbuster like *Ben Hur*, or *Waterworld* for that matter—results that will generate mixed feelings. So there are real limitations to the processes I describe.

But even acknowledging those limitations, it is fair to say that one of the most striking events to occur during our lifetimes is the transformation wrought by the Web, a transformation that is partly driven by the extraordinary explosion of nonproprietary creativity and sharing across digital networks. The cultural expectation that a web of expression and information will just *be* there—whatever subject we are discussing—is a fundamental one, the one that in some sense separates us from our children. With this as a background it is both bizarre and perverse that we choose to concentrate our policy making only on maintaining the business methods of the last century, only on the story line of the Internet Threat, only on the dangers that the technology poses to creativity (and it does pose some) and never on the benefits.

What would it mean to pay attention to the changes I have described? It would mean assessing the impact of rules on both proprietary and nonproprietary production. For example, if the introduction of a broad regime of software patents would render open source software development more difficult (because individual contributors cannot afford to do a patent search on every piece of code they contribute), then this should be reflected as a cost of software patents, to be balanced against whatever benefits the system brought. A method for encouraging innovation might, in fact, inhibit one form of it.

Paying attention to the last ten years means we need to realize that nonproprietary, distributed production is not the poor relation of traditional proprietary, hierarchically organized production. This is no hippy lovefest. It is the business method on which IBM has staked billions of dollars; the method of cultural production that generates much of the information each of us uses every day. It is just as deserving of respect and the solicitude of policy makers as the more familiar methods pursued by the film studios and proprietary software companies. Losses due to sharing that failed because of artificially erected legal barriers are every bit as real as losses that come about because of illicit copying. Yet our attention goes entirely to the latter.

The main thrust of the argument here is still firmly within the Jeffersonian, Scottish Enlightenment tradition. Jefferson does not wish to give the patent to Oliver Evans because he believes the invention will be (and has been) generated anyway without the granting of an intellectual property right and that there are sufficient information retrieval methods to have practical access to it. In this case, the information retrieval method is not Google. It is a polymath genius combing his library in Monticello for references to Persian irrigation methods. The “embarrassment” caused by the unnecessary patent is added expense and bureaucracy in agriculture and impediments to further innovators, not the undermining of open source software. But it is the same principle of cautious minimalism, the same belief that much innovation goes on without proprietary control and that intellectual property rights are the exception, not the rule. When Benjamin Franklin, a man who surely deserved patents under even the most stringent set of tests, chooses to forgo them because he has secured so much benefit from the contributions of others, he expresses Shirky’s norm nicely.

Indeed, Jefferson’s optimism depends partly on a view of information sharing that captures beautifully the attitudes of the generation that built the Web. The letter that I discussed in Chapter 2 was widely cited for precisely this reason. Remember these lines?

That ideas should freely spread from one to another over the globe, for the moral and mutual instruction of man, and improvement of his condition, seems to have been peculiarly and benevolently designed by nature, when she made them, like fire, expansible over all space, without lessening their density in any point, and like the air in which we breathe, move, and have our physical being, incapable of confinement or exclusive appropriation.

What could encapsulate better the process by which information spreads on a global network? What could more elegantly state the norms of the “information

wants to be free” generation? (Though those who quoted him conveniently omitted the portions of his analysis where he concedes that there are cases where intellectual property rights may be necessary and desirable.)

In some ways, then, the explosion of nonproprietary and, in many cases, noncommercial creativity and information sharing is simply the vindication of Jefferson’s comparison of ideas with “fire . . . expansible over all space.” The Web makes the simile a reality and puts an exclamation point at the end of the Jefferson Warning. All the more reason to pay attention to it. But the creative commons I described here goes further. It forces us to reconceptualize a form of life, a method of production, and a means of social organization that we used to relegate to the private world of informal sharing and collaboration. Denied a commons by bad intellectual property rules, we can sometimes build our own—which may in some ways do even more for us than the zone of free trade, free thought, and free action that Jefferson wished to protect.

Does all this mean that the Jefferson Warning is no longer necessary? Can we mitigate the negative effects of intellectual property expansion through a series of privately constructed commons? The answers to those questions are, respectively, “no” and “sometimes.” Think of the story of retrospectively extended copyright and orphan works. In many cases the problem with our intellectual property rights is that they create barriers to sharing—without producing an incentive in return—in ways that can never be solved through private agreement. Twentieth century culture will largely remain off-limits for digitization, reproduction, adaptation, and translation. No series of private contracts or licenses can fix the problem because the relevant parties are not in the room and might not agree if they were.

Even when the parties are available and agree to share, the benefits may not flow to all equally. Beset by a multitude of vague patents of questionable worth and uncertain scope, large information technology firms routinely create patent pools. IBM tosses in thousands of patents, so does Hewlett or Dell. Each agrees not to sue the other. This is great for the established companies; they can proceed without fear of legal action from the landmine patents that litter the technological landscape. As far as the participants are concerned, the patent pool is almost like the public domain—but a privatized public domain, a park that only residents may enter. But what about the start up company that does not have the thousands of patents necessary for entry? They are not in as happy a situation. The patent pool fixes the problem of poor patent quality and unclear scope—one that Jefferson was worrying about 200 years ago. But it fixes it only for the dominant firms, hurting competition in the process.

Attempts to form a commons may also backfire. The coordination problems are legion. There are difficulties of compatibility in licenses and the process, no matter how easy, still imposes transaction costs. Nevertheless, with all of these qualifications, the idea of the privately created commons is an important addition to the world view that Jefferson provided, a new tool in our attempt to craft a working system of innovation and culture. No one who looks at the Web can doubt the power of distributed, and frequently uncompensated, creativity in constructing remarkable reference works, operating systems, cultural conversations, even libraries of images and music. Some of that innovation happens largely outside of the world of intellectual property. Some of it happens in privately created areas of sharing that use property rights and open, sometimes even machine-readable, licenses to create a commons on which others can build. The world of creativity and its methods is wider than we had thought. That is one of the vital and exciting lessons the Internet teaches us; unfortunately, the only one our policy makers seem to hear is “cheaper copying means more piracy.”