


CHAPTER TWENTY-THREE
A Creative Commons?
Summary and Conclusion



We hope you have gleaned a number of things from this book. The first is the remarkable power of both rhetorical framing and moral and economic baselines in the debates we have been describing. If you understand them, it is like having a map of the entire area rather than tediously memorizing directions. Take trademark law. Earlier we summed up the two dominant visions of trademark.

One vision of the scope of a trademark confines it tightly to the semantic interaction between *this* good or service, *this* mark, *this* consumer and *this* manufacturer. Dove for soap does not infringe Dove for chocolate. The rationale for the right is to maintain stable meaning of the symbols that producers use and consumers rely on. The key *market* is the one the trademark owner is already in, with perhaps a small room for expansion to closely related markets (Levi designer jeans not Levi Jeep interiors). The key *person* whose confusion is relevant is the purchaser (and thus labels that explicitly disclaim connection to another producer are very strong evidence that consumers will not be likely to be confused). The key *moment* is the moment of sale. Confusion before the moment of sale (as with the person who clicks on a Google advertisement served up by her search for “Coach bags,” only to be shown Kate Spade bags that she ends up preferring to Coach) is irrelevant. Confusion *after* the moment of sale (as when the Chrysler 300 is seen on the street and mistaken for a Bentley, when the driver knows very well it is not) is irrelevant. Wherever the narrow reach of the right does not extend, competition—including competition built on deliberately copying non-trademarked features of another product—is to be welcomed.

The second, broader, vision of the scope of a trademark right views it partly as a device to avoid (current) consumer confusion and diminution of the utility of trademark symbols. In that it agrees with the narrower view. But it goes beyond that to see the trademark right as rooted in broader themes of unfair competition law, protecting acquired goodwill which can be leveraged into new markets whenever the producer wants, and preventing other producers from “reaping where they have not sown” even if the *consumer* is not at all confused. The right is no longer confined tightly to the semantic interaction between this consumer, this mark, this good and this producer. It is extended to cover possible *future* markets. It is as if, by having the mark, I have planted a semantic claim stake on the empty range next door. The relevant *moment* is expanded both before and after the point of sale, to cover initial interest confusion and post point of sale confusion. If the consumer was interested initially in my mark, I deserve to get that consumer, even if they come to prefer the goods of my rival. The person whose likelihood of confusion is relevant is not merely the actual consumer, but the

bystander and possible future purchaser. This vision of the right protects a larger swath of time, reaches a larger swath of markets and protects against (as a legal realist, Cohen would say ‘judicially creates’) “harms” that the first vision simply does not reach.

These two views of trademark law are of great importance in the doctrine. Compare genericide, *WalMart v. Samara* and the defense of fair or nominative use with *PETA*, initial interest confusion and anti-dilution law. Can you see the differing baselines of Pitney, Brandeis or Holmes here? Do you notice assumptions about whether positive externalities “naturally” (or efficiently) belong to the originator? See the way that the frames of competition and information-flow in the first view shade into the framing of property and entitlement in the second?

Our point is that these themes come up not just in trademark, but *throughout* intellectual property law. In some sense, we fight the same battles again and again—it would be nice to know that we are doing it, to understand the other side, to be a better advocate but also a better counselor and advisor.

Is computer code speech or property? Musical score or digital crowbar? Is an algorithm an invention running on a computer—and thus in the world of property? Or is it an idea, part of the great public domain of science, belonging in the world of free expression and interchange? Should we see the reverse engineering of a game-cartridge as the facilitating of competition by allowing interoperability? Or should we see it as a case of shamelessly copying, reaping where you have not sown? Is digitally restricting fair use just a property rights issue, like locking your diary in the safe, or is it an attempt to use copyright’s monopoly to restrict protected speech? Should trade secrets be protected only against acts that are illegal on other grounds, such as breach of confidentiality agreements? Or should the owner be protected even from legal actions, such as an airplane overflight that might gather valuable information, either because we think that it is immoral or because it leads to inefficient precautions? Yes, economic analysis can tell us much about these issues, but it too begins from implicit frames and baselines.

If one doubts all this, it is worth looking at the three baselines offered in *INS v. AP*, a case that is more than a century old, and seeing how they fit onto the emerging conflicts over generative AI, and in particular whether training AI models on copyrighted material should be a fair use. There you have the voice of Pitney (“AI companies are reaping where they have not sown, parasitically leaching away value from its true creators and owners!”). But we can also hear the more measured tones of Holmes, telling us that just because something is *valuable*, doesn’t mean there must be a property right over it. “Many exchangeable values may be destroyed intentionally without compensation.” If we are to make such a claim “some other ground must be found.” And there are definitely strains of Brandeis-thinking. The norm is freedom, property is the narrow exception to that norm. The preservation of a commons that can be freely copied is as vital to progress—here, the ability to build generative AI in the first place—as the incentives supplied by intellectual property rights.

True, there are changes in the way we argue. Partly in response to Holmes’ critique and partly because copyright law disdains labor-based arguments, the Pitney of today will be arguing that positive externalities must be internalized if efficient incentives are to result, and hypothesizing about efficient licensing markets rather than Biblical grain parables. Holmes had said that “some other ground [than mere uncompensated value] must be found” to justify an intellectual property right. Economic arguments supply that ground today. But those arguments are not monolithic. They also speak to the generative power, and frictionless efficiency, of commons-based production. Capitalism depends

on, indeed exalts, competition—in which freely copying valuable information produced by a study of one’s competitors is absolutely central. The transformative uses that fair use stresses are exactly such systems—a move to productive commons-based production in internet search, or book archiving (and perhaps AI), where formulaic application of the exclusive rights would have undermined the potential for disruptive innovation. This resonates deeply with the Brandeisian vision. Bottom line? Generative AI would have seemed strange to the judges in *INS*. The baselines, particularly those used in the popular debate, would be familiar. Understanding baselines is vital to future-proofing your legal skills. But it is also a fascinating guide to the powerful, and sometimes contradictory or conflicting, moral lenses through which our culture views both innovation and information.

The second theme was the effect of constitutional law—both the Intellectual Property clause and the First Amendment—on intellectual property. To summarize, we see a wild divergence between deference to Congress in copyright (*Eldred*, *Golan*) and—if *Graham* is still good law—a firmer hand in applying the Constitution’s limits to patent. We see that the Court is reluctant to strike down legislative provisions outright. It will do so if they seem flagrantly unconstitutional (*Tam*, *Brunetti*), though perhaps not if they look like traditional intellectual property rights (*Elster*). But we also see a willingness to be guided by constitutional purposes and the First Amendment’s strictures in *applying* intellectual property law. The constitution may exert its greatest sway on intellectual property *interstitially*.

The third theme was the adaptation of law to technological change. We had three main case studies: trademark law’s attempt to deal with domain names, copyright law’s attempts to deal with software, and patent law’s attempts to deal with genetic engineering on the one hand and the networked computer on the other. Generative AI offers a fourth case study that is still very much in progress. Each domain was presented with a *methodological* choice—does one formalistically apply the old rules to a reality that the technology has changed? Or does one focus on the purposive and utilitarian *goals* of the system and interpret the law accordingly—often using the old pigeon-holes of the law in new and surprising ways? The *PETA* decision in trademark, the *MAI* decision in copyright and—to a lesser extent—the CAFC’s disdain for “policy arguments” in the *Fisher* decision on utility standards in genetic patents represent the first tendency. By contrast, the *Playboy v. Welles* decision about nominative use and search engines in trademark, the *Lotus*, *Perfect 10* and *Google Books* cases in copyright and *some* aspects of the *Alice* case in patent represent the second. You should be able to come up with many others. We tried to present the arguments for both approaches. While admitting its difficulties, we are proponents of the second approach. You should make up your own mind. Which tendency will be exhibited in the decisions over AI? Our guess is “both.”

The fourth theme was that the limitations on and exceptions to intellectual property are as important as the rights themselves. “The holes matter as much as the cheese” to quote the discussion of *Sony*. In each field of intellectual property we saw the attempt to balance those things that should be in the realm of property and private right, and those that should be in the public domain, free for all to use. This balance is the recurrent *motif* of intellectual property policy. But how to *strike* that balance?

Again in each field, two tendencies manifested themselves. One sees strong rights as necessary and vital incentives. It views attempts to limit or constrain those rights as inherently problematic, a loss or “taking” from the property owner and a threat to innovation and economy. If X level of protection gives us Y level of innovation, then 2X will give us 2Y and so on. The norm is property. The best example of this was the discussion

of “The Internet Threat” in copyright policy and the idea that rights should vary inversely with the cost of copying. Costless copying will require nearly perfect—digitally enabled—control.

The other vision sees rights as “necessary evils”—Macaulay and Jefferson were its most articulate proponents. The right is a limited monopoly that has to be created to produce some social good: it should be held to the minimum scope, duration and extent necessary to achieve that goal. The norm is freedom, and the goal of intellectual property law is the crafting of limitations and exceptions to make sure the losses are as small as possible and the gains as great. The fair use cases presented many examples of this theme. How powerful is it? Well, if one looks at *Congressional* action—for example the repeated retrospective extensions of copyright—one has to say that it faces an uphill struggle. Will the fact that copyright law now directly affects citizens in a way it did not before change that fact? We do not know.

There is much that the book did not cover—though we wish we had time to do so. There has been very little here about the issues of distributive justice in intellectual property. Our patent system will never produce drugs that treat the diseases of the global poor. If one assesses value by ability and willingness to pay, as some but not all economic analysis does, obesity and hair loss treatments are more valuable than malaria vaccines. Most people find that result outrageous, correctly, but respond—wrongly in our view—by blaming pharmaceutical companies for it. We are the ones who set the system up that way. If we want drugs (and other inventions) that respond to the needs of the global poor, then other systems of incentives will be needed.

There *are* such systems—ranging from “prize funds” to cost-plus contracts to crowd-sourced innovation (though that last one will be of little help for the truly enormous problems of pharmaceutical development). Meanwhile the Access to Knowledge or A2K movement has stressed the role of liberal copyright policies in securing access to educational and cultural materials. These are deep and important issues and we invite you to read further on them.¹

But what of the future? In our final reading, we offer one last perspective on the changes that the internet may yet wreak on our assumptions about incentives, sharing and creativity.

James Boyle, A Creative Commons

The Public Domain, pp. 179–181, 183–200



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. Search for “David Byrne, My Fair Lady”

¹ See Knowledge Ecology International <http://keionline.org/> (visited July 20, 2018); Amy Kapczynski, *The Access to Knowledge Mobilization and the New Politics of Intellectual Property*, 117 YALE L.J. 804 (2008); Madhavi Sunder, *From Goods to a Good Life: Intellectual Property and Global Justice* (2012).

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, 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 ccMixter and you will find that musicians, famous and obscure, are inviting you to sample and remix their music.

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.”

...

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. . . .

[But actually] 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.

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. 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. 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. 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. . . .

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." 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?

. . .

Yochai Benkler and I would argue that these questions are fun to debate but ultimately irrelevant. 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. . . .

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. . . .

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. . . .

. . . 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.’” 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. This is at least a possible future, or part of a possible future, and one that we should not foreclose without thinking twice.

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,” 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?

Learning from the Sharing Economy

Lesson number one comes from 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. . . .

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. . . .

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 [twenty] 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.) . . .

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.

Questions:

1.) In Chapter 1, we read John Perry Barlow writing at the birth of the web and predicting its future effects on intellectual property. You were asked to assess what he got right and what he got wrong. You also read Cory Doctorow, from 2024. He claimed that law and policy—particularly in the realms of intellectual property and antitrust—need to be far more attentive to the economic and cultural dangers of “enshittification,” the result produced when expansive intellectual property rights and diminished antitrust enforcement meet a world of network effects. Twenty years from now, which parts of Boyle’s, Doctorow’s and Barlow’s arguments will seem naïve or dated and what parts—if any—correct?

2.) In this excerpt, Boyle argues that we should move away from policy made by anecdote and intuition, and instead move towards evidence-based policy that balances the contributions of both property and the public domain. In other words, we should make intellectual property policy the same way we make environmental rules or approve drugs through the FDA—based on criteria of efficacy and side effects. Wherever possible, we should outline clearly the result the new right is intended to have and then measure whether it creates it. We should use natural experiments—like that presented by Europe’s Database Directive—to assess the effects rights have. Is this approach realistic? What would Justice Breyer say? Justice Ginsburg?

3.) We have the widest access to culture we have ever had. It may be less than our archival technology would allow, particularly for works that are not commercially available, but it is still vastly greater than a person from 50 years ago would have considered possible. We also have continued, large-scale, illicit copying. The level of law-breaking may not be as high as it was during the Napster years. Cheap and convenient streaming has largely driven free but illegal downloads out of the music market, and much of the mainstream video market as well. But illicit copying is very much still a thing. So, one must concede that our world has high levels of cultural access, and continued illicit copying. Given this reality, why is Boyle complaining that our intellectual property law is *one-sided* and overly protectionist?
